Anisotropic Fast-Marching on cartesian grids using Lattice Basis Reduction

Jean-Marie Mirebeau*

September 24, 2012

Abstract

We introduce a modification of the Fast Marching Algorithm, which solves the anisotropic eikonal equation associated to an arbitrary continuous Riemannian metric $\mathcal{M}$, on a two or three dimensional box domain. The algorithm has a logarithmic complexity in the maximum anisotropy ratio $\kappa(\mathcal{M})$ of the Riemannian metric $\mathcal{M}$, which allows to handle extreme anisotropies for a reduced numerical cost. We establish that the output of the algorithm converges towards the viscosity solution of continuous problem, as the discretization step tends to zero. The algorithm is based on the computation at each grid point z of a reduced basis of the lattice $\mathbb{Z}^d$, with respect to the symmetric positive definite matrix $\mathcal{M}(z)$ encoding the desired anisotropy at this point.

Introduction

The eikonal equation, and its generalization the Hamilton-Jacobi equation, is a Partial Differential Equation (PDE) which describes an elementary front propagation model: the speed of the front depends only on the front position and orientation. This PDE is encountered in numerous applications, such as motion planning control problems [19], modeling of bio-medical phenomena [17], and image analysis [15]. It was also recently used in the context of medical image analysis [4] for extracting vessels in two dimensional projections or three dimensional scans of the human body, and for performing virtual endoscopies. This application requires to solve a highly anisotropic generalized eikonal equation with a high resolution on a cartesian grid, at a computational cost compatible with user interaction. It is one of the key motivations of this paper.

This paper is devoted to the construction and the study of a new algorithm, Fast Marching using Lattice Basis Reduction (FM-LBR), designed to solve the anisotropic eikonal equation associated to a given Riemannian metric $\mathcal{M}$, and which can handle large or even extreme anisotropies. The domain must be of dimension two or three, unless the metric has a special structure, see Point iii at the end of this introduction, and discretized on a cartesian grid. FM-LBR, as its name indicates, is a variant of the classical Fast Marching algorithm [19, 22], an efficient method for solving the eikonal equation when the metric is isotropic (proportional at each point to the identity matrix). Lattice Basis Reduction [14] is a concept from discrete mathematics, discussed in the first section of this paper, and used in the FM-LBR to produce local stencils for the discretization of the eikonal equation. Lattice Basis Reduction is involved here for the first time in the numerical analysis of a PDE to the knowledge of the author. We

*CNRS, University Paris Dauphine, UMR 7534, Laboratory CEREMADE, Paris, France.

describe the algorithm itself in the second section, where we also analyze its computational cost, and establish its consistency: the discrete approximations produced by this algorithm converge towards the viscosity solution of the continuous eikonal equation as the discretization step tends to zero. We present some numerical experiments in the third section, which confirm the small numerical cost of the algorithm and show an accuracy competitive in general, an remarkable in test cases related to the envisioned medical application.

We consider a positive integer d, the dimension, which is fixed throughout this paper. Let us mention that our algorithm only applies in dimension $d \in \{2, 3\}$, unless the problem of interest has a special structure, see Point iii at the end of this introduction. We denote by Ω the periodic unit box

$$\Omega := (\mathbb{R}/\mathbb{Z})^d. \tag{1}$$

The domain Ω coincides with the unit box $[-1/2,\, 1/2]^d$ equipped with periodic boundary conditions. We denote by $z + v \in \Omega$, the offset of a point $z \in \Omega$ by a vector $v \in \mathbb{R}^d$. The periodicity assumption is not essential, as shown in the numerical experiments and discussed in Point ii at the end of this introduction, but it simplifies the description and the proof of convergence of our algorithm. We consider a fixed Riemannian metric

$$\mathcal{M} \in C^0(\Omega, S_d^+)$$

where S_d^+ denotes the set of $d \times d$ symmetric positive definite matrices. For each $u \in \mathbb{R}^d$ and each $M \in S_d^+$ we denote $\|u\|_M := \sqrt{u^{\mathrm{T}} M u}$. Our objective is to compute (an approximation of) the *viscosity solution* $\mathrm{D} : \Omega \to \mathbb{R}_+$, see [11] and Definition 2.9, of the anisotropic eikonal equation

$$\begin{cases} \|\nabla \mathrm{D}(z)\|_{\mathcal{M}(z)^{-1}} = 1 & \text{for almost every } z \in \Omega \setminus \{0\}, \\ \mathrm{D}(0) = 0. \end{cases} \tag{2}$$

The only element of general theory used in this paper is the *uniqueness* of the viscosity solution to the above problem. Another point of view on this problem is provided by the following characterization [11]: for each $z \in \Omega$, the quantity $\mathrm{D}(z)$ is the length of the shortest path γ joining z to the origin. More precisely the solution of the eikonal equation is

$$\mathrm{D}(z) = \mathrm{D}(z, 0), \tag{3}$$

where $\mathrm{D}(\cdot,\cdot)$ denotes the Riemannian distance, defined for all $x, y \in \Omega$ by

$$\mathrm{D}(x, y) \;:=\; \inf\{\mathrm{length}(\gamma);\, \gamma \in C^1([0,1], \Omega),\, \gamma(0) = x,\, \gamma(1) = y\}, \tag{4}$$

$$\mathrm{length}(\gamma) \;:=\; \int_0^1 \|\gamma'(t)\|_{\mathcal{M}(\gamma(t))} dt. \tag{5}$$

We define the anisotropy ratio $\kappa(M)$ of a matrix $M \in S_d^+$, and the maximum anisotropy $\kappa(\mathcal{M})$ of the Riemannian metric $\mathcal{M}$, as follows (denoting by $\|\cdot\|$ the standard euclidean norm)

$$\kappa(M) := \max_{\|u\| = \|v\| = 1} \frac{\|u\|_M}{\|v\|_M}, \quad \kappa(\mathcal{M}) := \max_{z \in \Omega} \kappa(\mathcal{M}(z)). \tag{6}$$

Note that $\kappa(M) = \sqrt{\|M\|\|M^{-1}\|}$. The test cases presented in §3 cover both moderate anisotropy $\kappa(\mathcal{M}) \simeq 5$, and "extreme" anisotropy $\kappa(\mathcal{M}) \approx 100$, which is relevant in applications to image segmentation and structure extraction [4].

The periodic box Ω is discretized on a cartesian grid Ω_n of step size $1/n$

$$\Omega_n := \{0,\, 1/n,\, 2/n, \cdots,\, (n-1)/n\}^d \subset \Omega.$$

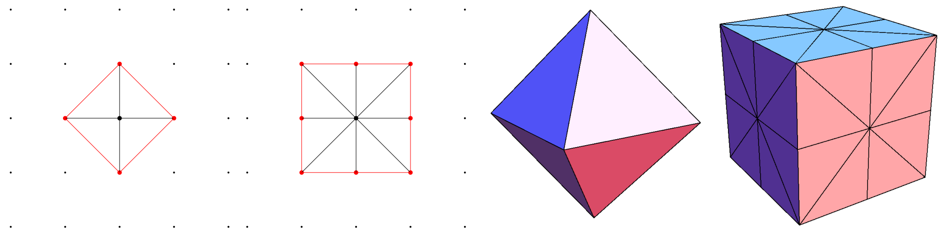

Figure 1: Some classical stencils used in the discretization of two dimensional (left) or three dimensional (right) eikonal equations, associated to a Riemannian metric $\mathcal{M}$. A Dijkstra inspired method, using one of these stencils at a point $z \in \Omega_n$, will be consistent if either $\mathcal{M}(z)$ is diagonal, or $\kappa(\mathcal{M}(z)) \leq \kappa_0$ where $\kappa_0 = 1,\ 1+\sqrt{2},\ 1,\ (\sqrt{3}+1)/2$ (from left to right). See Proposition 1.2.

Denoting by $N := n^d = \#(\Omega_n)$ the number of grid points, our algorithm has complexity

$$\mathcal{O}(N \ln N + N \ln \kappa(\mathcal{M})), \tag{7}$$

where the constant behind the $\mathcal{O}(\cdot)$ notation only depends on the dimension d. In practical applications one has $N \gtrsim 10\,000$ and $\kappa(\mathcal{M}) \lesssim 100$, hence the first term dominates in (7), and the complexity of our algorithm is only $\mathcal{O}(N \ln N)$.

Let $z \in \Omega$ and let $V \subset \Omega$ be a compact neighborhood of z, not containing the origin in its interior. It follows from the interpretation of D as a path-length distance (3) that

$$\mathrm{D}(z) = \min_{z' \in \partial V} \mathrm{D}(z, z') + \mathrm{D}(z'). \tag{8}$$

Our algorithm, as well as most alternative solvers of the eikonal equation, discretizes this PDE using an approximation of the right hand side of (8): the Hopf-Lax update operator [19, 5], earliest references in [9, 6]. Equation (8) is then reinterpreted as a fixed point problem. More precisely, consider a fixed $n \geq 1$ and assume that a small simply connected neighborhood $V_n(z)$ of each $z \in \Omega_n$ has been constructed *under the form of a simplicial mesh* (a triangulation if $d = 2$, a tetrahedral mesh if $d = 3$, see Figure 1), with all its vertices in Ω_n.

We denote $\overline{\mathbb{R}}_+ := \mathbb{R}_+ \cup \{+\infty\}$, equipped with the topology of a compact segment, and we adopt the convention $0 \times \infty = 0$. For any discrete map $\mathrm{d} : \Omega_n \to \overline{\mathbb{R}}_+$, the Hopf-Lax update operator is defined by

$$\Lambda_n(\mathrm{d}, z) := \min_{z' \in \partial V} \|z' - z\|_{\mathcal{M}(z)} + \mathrm{I}_V\, \mathrm{d}(z'), \tag{9}$$

where $z \in \Omega_n$ is arbitrary, has neighborhood $V := V_n(z)$, and I_V denotes the piecewise linear interpolation operator on this simplicial mesh. The difference $z' - z$ is well defined in (9) because z and z' both belong to the small and simply connected set $V \subset \Omega$, the stencil of z, which can be identified to a subset of $\mathbb{R}^d$ (this convention will not be recalled systematically in the rest of the paper). The Hopf-Lax update operator Λ_n approximates the distance $\mathrm{D}(z, z')$ in (8) by the local norm $\|z' - z\|_{\mathcal{M}(z)}$, for $z' \in \partial V$, and the value $\mathrm{D}(z')$ by linear interpolation of d on the mesh V. Note that $\Lambda_n(\mathrm{d}, z)$ depends on $\mathrm{d}(z')$ for all vertices $z' \in \Omega_n \cap \partial V$, on the intersection of the discretization grid with the stencil boundary.

The discrete approximation $\mathrm{d}_n : \Omega_n \to \mathbb{R}_+$, of the continuous solution D of the eikonal equation, is obtained as the solution of the N-dimensional fixed point problem

$$\begin{cases} \Lambda_n(\mathrm{d}_n, z) = \mathrm{d}_n(z) & \text{for all } z \in \Omega_n \setminus \{0\}, \\ \mathrm{d}_n(0) = 0. \end{cases} \tag{10}$$

Numerical solvers of the eikonal equation differ by (i) the construction of the stencils $V_n(z)$, $z \in \Omega_n$, and (ii) the approach used to solve the system (10) which is inspired by the algorithms of Bellmann-Ford or of Dijkstra used in graph theory. The algorithm presented in this paper, FM-LBR, belongs to the category of Dijkstra inspired algorithms with static stencils, and among these is the first one to guarantee a uniform upper bound on the stencil cardinality.

- **Bellman-Ford inspired algorithms.** The discrete fixed point problem (10) is solved via Gauss-Seidel iteration: the replacement rule $\mathrm{d}(z_k) \leftarrow \Lambda_n(\mathrm{d}, z_k)$ is applied for $k = 0, 1, 2, ...$ to a mutable map $\delta : \Omega_n \to \mathbb{R}$, until a convergence criterion is met. In the fast sweeping methods, see [21] and references therein, the sequence of points $(z_k)_{k\geq 0}$ enumerates repeatedly the lines and the columns of Ω_n. Alternatively this sequence is obtained via a priority queue in the Adaptive Gauss-Seidel Iteration (AGSI) of Bornemann and Rasch [5]. The stencil $V(z)$ of a point $z \in \Omega_n$ is usually the offset by z of a fixed stencil V given at the origin, such as those illustrated on Figure 1.

  Fast sweeping methods have $\mathcal{O}(\lambda(\mathcal{M})N)$ complexity when the metric $\mathcal{M}$ is *isotropic* (proportional to the identity at each point), but this result does not extend to anisotropic Riemannian metrics, see [26] for the proof and the expression of $\lambda(\mathcal{M})$. The AGSI has complexity $\mathcal{O}(\mu(\mathcal{M})N^{1+\frac{1}{d}})$, for arbitrary anisotropic Riemannian metrics, where $\mu(\mathcal{M})$ is a non explicit constant which depends on global geometrical features of the metric [5]. The AGSI is a popular, simple, and quite efficient method, which is included for comparison in our numerical tests.

- **Dijkstra inspired algorithms.** The system (10) is solved in a single pass, non-iteratively, using an ordering of Ω_n determined at run-time. This is possible provided the Hopf-Lax update operator satisfies the so-called "causality property", see Lemma 2.3, which can be ensured if the stencil $V(z)$ of each $z \in \Omega_n$ satisfies some geometrical properties depending on $\mathcal{M}(z)$, see Definition 1.1. The different Dijkstra inspired methods are characterized by the construction of the stencils $V(z)$, which depend on $\mathcal{M}(z)$, in contrast with Bellman-Ford inspired methods which are characterized by the choice of the sequence $(z_k)_{k\geq 0}$. Solving the system (10) with a Dijkstra inspired algorithm has complexity $\mathcal{O}(\mu(\mathcal{M})N \ln N)$, where $\mu(\mathcal{M})$ is an upper bound for the cardinality of the stencils (the number of simplices they are built of).

  In the Ordered Upwind Method (OUM) of Sethian and Vladimirsky [19, 25], the stencils are constructed at run-time; their cardinality is bounded by $\mathcal{O}(\kappa(\mathcal{M})^d)$ and drops to $\mathcal{O}(\kappa(\mathcal{M})^{d-1})$ as $N \to \infty$. In contrast, the stencils are constructed during a preprocessing step and then static in the Monotone Acceptance Ordered Upwind Method (MAOUM) of Alton and Mitchell [2]; their cardinality is bounded by $\mathcal{O}(\kappa(\mathcal{M})^d)$. The FM-LBR introduced in the present work uses an approach similar to the MAOUM, except that the cardinality of the stencils is $\mathcal{O}(1)$, fully independent of the Riemannian metric $\mathcal{M}$. The complexity estimates are thus $\mathcal{O}(\kappa(M)^d N \ln N)$ for the OUM and the MAOUM (asymptotically $\mathcal{O}(\kappa(M)^{d-1} N \ln N)$ for the OUM), and $\mathcal{O}(N \ln N + N \ln \kappa(\mathcal{M}))$ for our approach, FM-LBR, where the second term in the complexity accounts for the stencil construction.

The above mentioned algorithms are consistent for the anisotropic eikonal equation associated to an arbitrary continuous Riemannian metrics $\mathcal{M} : \Omega \to S_d^+$, in the sense that the discrete output d_n of the algorithm converges to the viscosity solution D of the continuous problem as $n \to \infty$. Some more specialized variants of the fast marching algorithm are only consistent for a restricted set of metrics, but can be executed nonetheless with an arbitrary anisotropic

metric $\mathcal{M}$ (In that case the discrete system (10) may not be solved, and the numerical results are variable, see §3). For instance the original fast marching algorithm [22] is consistent if $\mathcal{M}(z)$ is proportional to the identity matrix for each $z \in \Omega$, and more generally if $\mathcal{M}(z)$ is a diagonal matrix. In addition to these cases of isotropy and axis-aligned anisotropy, some variants [7] are also consistent if $\kappa(\mathcal{M}) \leq \kappa_0$, where κ_0 is a given bound, see Figure 1 for some examples in two and three dimensions. Our numerical experiments include for comparison one of these methods: Fast Marching using the 8 point stencil (FM-8, center left stencil on Figure 1), which is popular in applications [4] thanks to its short computation times and despite the lack of convergence guarantee for arbitrary metrics. Depending on the implementation [16, 19, 22], involving either a sorted list or a bucket sort, these methods have complexity $\mathcal{O}(N \ln N)$ or $\mathcal{O}(\Upsilon(\mathcal{M})N)$, where

$$\Upsilon(\mathcal{M}) := \sqrt{\max_{z\in\Omega} \|\mathcal{M}(z)\| \max_{z'\in\Omega} \|\mathcal{M}(z')^{-1}\|}.$$

In the applications for which our method is intended, one typically has $\ln(N) \lesssim \kappa(\mathcal{M}) \leq \Upsilon(\mathcal{M}) \ll N$, in such way that the complexity (7) of the proposed method is comparable to $\mathcal{O}(N \ln N)$ and smaller than $\mathcal{O}(\Upsilon(\mathcal{M})N)$. In summary the algorithm proposed in this paper has a complexity $\mathcal{O}(N \ln N + N \ln \kappa(\mathcal{M}))$ which is significantly less than general methods such as the AGSI, the OUM and the MAOUM, which are consistent for arbitrary Riemannian metrics, and no more than the complexity of more specialized methods, such as the original fast marching algorithm.

Our complexity analysis guarantees short and predictable run-times, which is desirable in applications involving user interaction, e.g. image processing. It does not answer however the question of accuracy.

Remark 1 (Accuracy and stencil construction)**.** *The above compared solvers of the eikonal equation, the AGSI, the OUM, the MAOUM and our method the FM-LBR, involve different stencils and thus a different Hopf-Lax update operator Λ_n. The discrete fixed point problem (10) therefore depends on the method, and so does its solution d_n, which impacts the accuracy of the scheme. In the special case of a* constant *Riemannian metric, the theoretical error analysis presented in [12], shows that the anisotropy ratio $\kappa(\mathcal{M})$ does not impact the accuracy of FM-LBR, in an average sense over grid orientations. In contrast one expects the numerical error of alternative methods to grow proportionally to $\kappa(\mathcal{M})$, as discussed in [19] and suggested in [12] by numerical experiments for the AGSI. This error analysis does not extend to general Riemannian metrics, and we do not claim that the proposed algorithm outperforms its alternatives in terms of accuracy in all applications. The numerical experiments presented in section §3 show nevertheless that FM-LBR is competitive in this regard.*

Let us emphasize that the efficiency of the proposed algorithm, the FM-LBR, comes at the price of its specialization. We review below its limitations:

i (Finsler metrics) FM-LBR only applies to the anisotropic eikonal equation associated to a *Riemannian* metric. The underlying Riemannian structure plays an important role in our approach, and our algorithm therefore cannot handle more general Hamilton-Jacobi equations, contrary to the AGSI [5], the OUM [19] and the MAOUM [2] mentioned above. These approaches allow Finsler metrics in addition to Riemannian metrics, in other words the local euclidean norm $\|\cdot\|_{\mathcal{M}(z)}$ may be replaced for each $z \in \Omega$ with an arbitrary asymmetric norm $|\cdot|_z \in C^0(\mathbb{R}^d, \mathbb{R}+)$.

  Constructing static stencils suitable for Dijkstra inspired algorithms, in the sense that they guarantee the causality property as in Lemma 2.3 (the stencil is said to be *causal*), is an

active subject of research in the case of Finsler metrics. A characterization was obtained in [23], and a construction proposed in [2, 13]. It was also observed in [18, 19, 1] that the canonical stencil (Figure 1, left) is causal for a large variety of axis-aligned Finsler metrics, in such way that the original fast marching algorithm can be applied.

ii (Domain discretization) FM-LBR requires a domain discretized on a cartesian grid. Our algorithm therefore does not apply to domains provided under the form of general unstructured meshes, a more difficult setting which has attracted an important research effort [3, 5, 8, 10, 19]. At the price of a higher technicality of the proof, the specific periodic domain (1) could however be replaced with a more general smooth bounded domain $\Omega \subset \mathbb{R}^d$, still discretized on a grid. In that case the PDE (2) is replaced with

$$\begin{cases} \|\nabla \mathrm{D}(z)\|_{\mathcal{M}(z)^{-1}} = 1 & \text{for almost every } z \in \Omega, \\ \mathrm{D}(z) = f(z) & \text{for all } z \in \partial\Omega. \end{cases} \tag{11}$$

where the boundary data f needs to satisfy [11] the Lipschitz regularity condition $|f(z) - f(z')| \leq \text{length}(\gamma)$ for any path $\gamma \in C^1([0,1], \overline{\Omega})$ joining two points $z, z' \in \partial\Omega$. Note that using the FM-LBR in this context will require the extension of the boundary data to a ghost layer covering $\partial\Omega$, since this algorithm uses large stencils, of euclidean diameter $\mathcal{O}(\kappa(\mathcal{M})/n)$ instead of e.g. $\mathcal{O}(1/n)$ for the AGSI.

iii (Dimension) FM-LBR applies to arbitrary continuous Riemannian metrics on a two or three dimensional domain (an extension to four dimensional domains is presented in [12]). The algorithm can be extended to higher dimension if the Riemannian metric $\mathcal{M}$ has a specific diagonal block structure, with blocks of size 1, 2 or 3, see §2. Such block diagonal structures are not uncommon in the context of medical imaging, see [4]. They are inherited from the cartesian product structure of the fast marching domain: $\Omega = \Omega_0 \times \Omega_1$, where Ω_0 is a physical domain of dimension ≤ 3, and Ω_1 is an abstract parameter domain of dimension ≤ 2.

We introduce and study in §1 the notion of M-reduced mesh, where $M \in S_d^+$ is a symmetric positive definite matrix. The construction of M-reduced meshes of bounded cardinality, using Lattice Basis Reduction, is the main originality of this paper and the key of the small complexity of our algorithm the FM-LBR, which uses them as stencils. Following a more classical approach, we describe in §2 the FM-LBR as a variant of the Fast Marching algorithm using this specific stencil construction, and we establish the related convergence result. We finally we present some numerical experiments in §3.

Remark 2 (Computing distances on a surface)**.** *Consider a smooth surface $\mathcal{S} \subset \mathbb{R}^3$ (or more generally a smooth embedded manifold $\mathcal{S} \subset \mathbb{R}^d$), equipped with Riemannian metric induced by the euclidean metric on $\mathbb{R}^3$. In order to compute distances on $\mathcal{S}$, a first possibility is to solve an* isotropic *eikonal equation on $\mathcal{S}$, by applying the OUM on a triangulated mesh of $\mathcal{S}$ [20], or the original fast marching algorithm if this triangulation only contains acute triangles. A second possibility is to use one or several local charts $\varphi : \Omega \to \mathcal{S}$ and to solve an* anisotropic *eikonal equation on Ω; grid discretisations are allowed since $\Omega \subset \mathbb{R}^2$, hence the FM-LBR can be applied, as well as numerous alternatives. In practice the data usually dictates the choice between these approaches, since changing from one representation to the other (triangulation or charts) is a non-trivial and computationally intensive procedure.*

1 Reduced meshes and bases

We introduce in this section the notion of M-reduced meshes, where $M \in S_d^+$ is a given symmetric positive definite matrix. Anticipating on §2, we motivate their study by their use as static stencils for Dijkstra inspired eikonal solvers, such as the FM-LBR. Consider an Riemannian metric $\mathcal{M} \in C^0(\Omega, S_d^+)$, and assume that for each z in the grid Ω_n a $\mathcal{M}(z)$-reduced mesh $\mathcal{T}(z)$ has been constructed. Denote by $V_n(z)$ the stencil obtained by rescaling and offsetting the mesh $\mathcal{T}(z)$: with obvious notations

$$V_n(z) = z + \frac{1}{n}\mathcal{T}(z). \tag{12}$$

We show in §2 that the discrete fixed point system (10), obtained with these stencils, can be solved in a single pass, and we establish a convergence result towards the solution of the continuous eikonal equation, in Theorem 2.2.

A simplex $T \subset \mathbb{R}^d$ is the convex hull of $d+1$ points $v_0, \cdots, v_d \in \mathbb{R}^d$, not lying on a common hyperplane, which are called the vertices of T. A mesh is a finite collection $\mathcal{T}$ of simplices which satisfy the following conformity condition: for all $S, T \in \mathcal{T}$ the intersection $S \cap T$ is the convex hull of the common vertices of S and T.

Definition 1.1. *A M-reduced mesh is a mesh $\mathcal{T}$ which satisfies the following properties.*

(a) The union of the elements of $\mathcal{T}$ is a neighborhood of the origin

(b) The vertices of each simplex $T \in \mathcal{T}$ lie on the lattice $\mathbb{Z}^d$, and T has volume $1/d!$.

(c) For each $T \in \mathcal{T}$, one of the vertices of T is the origin 0, and the others denoted by $v_1, \cdots, v_d$ satisfy for all $1 \leq i \leq j \leq d$

$$v_i^{\mathrm{T}} M v_j \geq 0. \tag{13}$$

Heuristically Point (a) of Definition 1.1 ensures that the numerical information is propagated in all directions in the discrete fixed point system (10). Point (b) implies that the non-zero vertices of any simplex $T \in \mathcal{T}$ form a basis of $\mathbb{Z}^d$, see Definition 1.4, a property later used to show that the information does not "fly over" a subset of the grid Ω_n. Point (c) is an acuteness condition related to the celebrated Causality property, see [19], which is a prerequisite for the Fast Marching algorithm and other one-pass solvers of eikonal or Hamilton-Jacobi equations.

The next proposition gives a simple criterion to show that a given mesh $\mathcal{T}$ is M-reduced for all $M \in S_d^+$ of sufficiently small anisotropy $\kappa(M)$.

Proposition 1.2. *Let $\mathcal{T}$ be a d-dimensional mesh which satisfies the requirements (a) and (b) of Definition 1.1. Let*

$$\kappa(\mathcal{T}) := \sqrt{\frac{1+\gamma(\mathcal{T})}{1-\gamma(\mathcal{T})}}, \quad \textit{where} \quad \gamma(\mathcal{T}) := \min_{T,(u,v)} \frac{u^{\mathrm{T}} v}{\|u\|\|v\|}, \tag{14}$$

and where the minimum in $\gamma(\mathcal{T})$ is taken a non-zero vertices u, v of a common simplex $T \in \mathcal{T}$. The mesh $\mathcal{T}$ is M-reduced for any $M \in S_d^+$ such that $\kappa(M) \leq \kappa(\mathcal{T})$.

Proof. Let u, v be two non-zero vertices of a common simplex $T \in \mathcal{T}$, and let $M \in S_d^+$. Let $u' := u/\|u\|$ and let $v' := v/\|v\|$. By construction we have

$$\|u'+v'\|^2 = 2(1+u'^{\mathrm{T}} v') \geq 2(1+\gamma(\mathcal{T})), \qquad \|u'-v'\|^2 = 2(1-u'^{\mathrm{T}} v') \leq 2(1-\gamma(\mathcal{T})).$$

Let us assume for contradiction that $u^{\mathrm{T}}Mv < 0$, which implies that $\|u'+v'\|_M < \|u'-v'\|_M$. Observing that

$$\kappa(M)^2 = \|M\|\|M^{-1}\| \geq \frac{\|u'-v'\|_M^2}{\|u'-v'\|^2}\frac{\|u'+v'\|^2}{\|u'+v'\|_M^2} > \frac{1+\gamma(\mathcal{T})}{1-\gamma(\mathcal{T})},$$

we obtain that $\kappa(M) > \kappa(\mathcal{T})$, which concludes the proof of this proposition. □

The meshes illustrated on Figure 1 are easily checked to satisfy requirements (a) and (b) of Definition 1.1. Requirement (c) clearly holds for diagonal matrices. Evaluating (14), we obtain $\kappa(\mathcal{T}) = 1,\ 1+\sqrt{2},\ 1,\ (\sqrt{3}+1)/2$ (from left to right), hence (c) also holds for matrices $M \in S_d^+$ of anisotropy $\kappa(M)$ within this bound. These meshes are therefore $\mathcal{M}(z)$-reduced under the conditions described in the caption of Figure 1, hence the resulting system of equations (10) can be solved in one pass by a Dijkstra inspired method.

Remark 1.3 (Block diagonal matrices). *Let d_1, d_2 be positive integers, let $M_1 \in S_{d_1}^+$, and let $M_2 \in S_{d_2}^+$. Let $\mathcal{T}_1$ be a M_1-reduced mesh, and let $\mathcal{T}_2$ a M_2-reduced mesh. Let $d := d_1 + d_2$ and let $\mathcal{T}$ be the d-dimensional mesh defined as follows: for any $T_1 \in \mathcal{T}_1$ of vertices $0 = u_0^1, \cdots u_{d_1}^1$, and any $T_2 \in \mathcal{T}_2$ of vertices $0 = u_0^2, \cdots, u_{d_2}^2$, the d-dimensional simplex T of vertices*

$$(0,0), (u_1^1, 0), \cdots, (u_{d_1}^1, 0), (0, u_1^2), \cdots, (0, u_{d_2}^2)$$

belongs to $\mathcal{T}$. Then one easily checks that $\mathcal{T}$ is a M-reduced mesh, where $M \in S_d^+$ denotes the matrix of diagonal blocks M_1 and M_2.

We introduce in subsection §1.1 the algebraic notion of M-reduced basis of $\mathbb{Z}^d$, where $M \in S_d^+$ and $1 \leq d \leq 4$. We show that the collection of vertices of a M-reduced mesh contains a M-reduced basis, a property later used in the analysis of our algorithm. We next give in §1.2 an explicit construction of a M-reduced mesh of bounded cardinality, for each $M \in S_d^+$, $d \in \{2,3\}$.

1.1 Bases of the lattice $\mathbb{Z}^d$

The results of this section fall in the framework of low-dimensional lattice basis reduction. We refer to [14] and references therein for an introduction to this rich theory, from which we use only one result: Theorem 1.5 stated below.

We denote by $u_1\mathbb{Z} + \cdots + u_k\mathbb{Z}$ the sub-lattice of $\mathbb{Z}^d$ generated by $u_1, \cdots, u_k \in \mathbb{Z}^d$:

$$u_1\mathbb{Z} + \cdots + u_k\mathbb{Z} := \{u_1 z_1 + \cdots + u_k z_k;\ z_1, \cdots, z_k \in \mathbb{Z}\}.$$

If $k = 0$ then the above sum equals $\{0\}$ by convention.

Definition 1.4. *A basis of $\mathbb{Z}^d$ is a d-plet $(u_1, \cdots, u_d)$ of elements of $\mathbb{Z}^d$ such that*

$$|\det(u_1, \cdots, u_d)| = 1. \tag{15}$$

Assume that $1 \leq d \leq 4$. A M-reduced basis of $\mathbb{Z}^d$, where $M \in S_d^+$, is a basis $(u_1, \cdots, u_d)$ of $\mathbb{Z}^d$ which satisfies for all $1 \leq k \leq d$

$$u_k \in \operatorname{argmin}\{\|z\|_M;\ z \in \mathbb{Z}^d \setminus (u_1\mathbb{Z} + \cdots + u_{k-1}\mathbb{Z})\}. \tag{16}$$

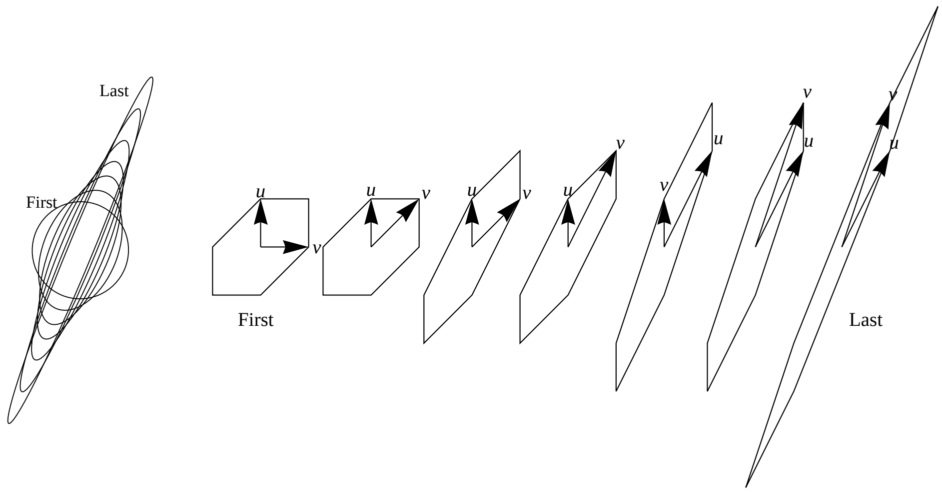

Figure 2: The unit sphere $\{u;\ \|u\|_M = 1\}$, a M-reduced basis (u, v), and the boundary of a M-reduced mesh constructed as in Proposition 1.9, for some $M \in S_2^+$ of anisotropy ratio $\kappa(M)$ ranging from 1 to 15, and eigenvector $(\cos(3\pi/8), \sin(3\pi/8))$ associated to the small eigenvalue.

For instance the canonical basis $(e_1, \cdots, e_d)$ of $\mathbb{R}^d$, is also a basis of $\mathbb{Z}^d$. If $M \in S_d^+$ is a diagonal matrix of coefficients $(\lambda_1, \cdots, \lambda_d)$, and if $0 < \lambda_{\sigma(1)} \leq \cdots \leq \lambda_{\sigma(d)}$ for some permutation σ, then $(e_{\sigma(1)}, \cdots e_{\sigma(d)})$ is a M-reduced basis of $\mathbb{Z}^d$. See Figures 2, 3, 4 and 5 for some examples of M-reduced bases associated to non-diagonal matrices M.

A M-reduced basis of $\mathbb{Z}^d$ contains by definition *small* vectors with respect to the norm $\|\cdot\|_M$: the smallest linearly independent ones with integer coordinates. As a result, the construction of M-reduced meshes based on M-reduced bases, presented in the next subsection, will produce *small* meshes, again with respect to the norm $\|\cdot\|_M$. As a result, the discretization of the FM-LBR uses *small stencils* $V_n(z)$ (12) at grid points $z \in \Omega_n$, in the sense of the local norm $\|\cdot\|_{\mathcal{M}}(z)$. Small stencils often the key of good accuracy in the discretization of PDEs. This observation is the starting point of [12], an error analysis of the FM-LBR, in the case of constant Riemannian metrics.

The above definition of a M-reduced basis of $\mathbb{Z}^d$ is adequate only up to dimension 4, since in dimension $d \geq 5$ there exists matrices $M \in S_d^+$ such that no basis of $\mathbb{Z}^d$ satisfies the relations (16), see [14] (these relations state that $\|u_i\|_M$ equals the i-th Minkowski's minimum $\lambda_i(M)$). The proper generalization of Definition 1.4 in dimension $d \geq 5$, is Minkowski's reduction [14].

Theorem 1.5 (Nguyen, Stelhé, 2009)**.** *There exists an algorithm which, given a matrix $M \in S_d^+$ as input, $1 \leq d \leq 4$, produces a M-reduced basis of $\mathbb{Z}^d$ and has the numerical cost $\mathcal{O}(1+\ln\kappa(M))$.*

Proof. The proof is contained in [14], and we only point out here the precise reference within the paper and the slight differences in notations. The algorithm described in [14] takes as input a basis $(b_1, \cdots, b_d)$ (here: the canonical basis of $\mathbb{R}^d$) of a lattice L (here: $\mathbb{Z}^d$), and its Gram matrix with respect to some scalar product (here: the Gram matrix is M). The algorithm outputs a greedy reduced basis of the lattice L, a notion which coincides with Minkowski's reduction if $d \leq 4$ (Lemma 4.3.2 in [14]), which itself coincides with Definition 1.4 if $d \leq 4$.

The main loop of the iterative algorithm is executed at most the following number of times (Theorem 6.0.5 in [14]):

$$\mathcal{O}\left(1 + \ln \max_{1 \leq i \leq d} \|b_i\|_M - \ln \min_{u \in L} \|u\|_M\right),$$

hence $\mathcal{O}(1+\ln\|M\|^{\frac{1}{2}}-\ln\|M^{-1}\|^{-\frac{1}{2}})=\mathcal{O}(1+\ln\kappa(M))$ times in our setting. The complexity of each of these iterations is dominated by a closest vector search, described in Theorem 5.0.4 in [14], which consists of the inversion of a $k\times k$ Gram matrix, where $1\leq k\leq d-1$, and a $\mathcal{O}(1)$ exhaustive search. In terms of elementary operations $(+,-,\times,/)$ among reals, each iteration of this algorithm thus has cost $\mathcal{O}(1)$, and the overall cost is the number of iterations $\mathcal{O}(1+\ln\kappa(M))$.

Note that an important part of the discussion in [14] is devoted to the special case where the vectors $(b_1,\cdots,b_d)$ have *large integer* coefficients, the Gram matrix is computed with respect to the standard euclidean scalar product, and the complexity of an elementary operation $(+,-,\times,/)$ among integers is not $\mathcal{O}(1)$ but depends on the size of these integers. This more subtle notion of complexity, named bit complexity, is not relevant in our setting. □

In dimension $d=2$, the algorithm mentioned un Theorem 1.5 mimicks the search for the greatest common divisor of two integers, and is often referred to as Gauss's algorithm [14]. This algorithm uses only a pair (u,v) of (mutable) variables in $\mathbb{Z}^2$, initialized as the canonical basis of $\mathbb{R}^2$. The pair (u,v) becomes a M-reduced basis at the end of the following loop, which takes at most $\mathcal{O}(\ln\kappa(M))$ iterations. Round denotes rounding to a closest integer.

$$\textbf{Do } (u,\ v) ::= (v,\ u-\mathrm{Round}(u^{\mathrm{T}}Mv/\|v\|_M^2)\,v)\textbf{, while } \|u\|_M>\|v\|_M\textbf{.}$$

Proposition 1.6. *Assume that $1\leq d\leq 4$. Let $M\in S_d^+$ and let $(u_1,\cdots,u_d)$ be a M-reduced basis of $\mathbb{Z}^d$. Then for all $1\leq i\leq d$*

$$\begin{aligned}\|u_i\| &\leq \kappa(M), &(17)\\ \|u_i\|_M &\leq \kappa(M)\|u_1\|_M. &(18)\end{aligned}$$

For any integer combination z of the elements of the basis distinct from u_i, in other words $z=\alpha_1u_1+\cdots\alpha_{i-1}u_{i-1}+\alpha_{i+1}u_{i+1}+\alpha_du_d$, where $\alpha_1,\cdots,\alpha_{i-1},\alpha_{i+1},\cdots,\alpha_d\in\mathbb{Z}$, one has

$$2|u_i^{\mathrm{T}}Mz|\leq\|z\|_M^2. \tag{19}$$

Proof. We consider a fixed $1\leq i\leq d$, and we claim that

$$\|M^{-1}\|^{-\frac{1}{2}}\leq\|u_1\|_M\leq\|u_i\|_M\leq\|M\|^{\frac{1}{2}}. \tag{20}$$

The left inequality follows from $1\leq\|u_1\|\leq\|M^{-\frac{1}{2}}\|\|u_1\|_M$. The central inequality follows from the fact that u_1 minimizes the norm $\|\cdot\|_M$ among all elements of $\mathbb{Z}^d\setminus\{0\}$, see Definition 1.4. Denoting by $(e_1,\cdots,e_d)$ the canonical basis of $\mathbb{R}^d$, we observe comparing dimensions that there exists $1\leq j\leq d$ such that $e_j\notin u_1\mathbb{Z}+\cdots+\cdots u_{i-1}\mathbb{Z}$. Therefore $\|u_i\|_M\leq\|e_j\|_M\leq\|M\|^{\frac{1}{2}}$ using Definition 1.4, which establishes (20), hence also (18). We obtain (17) combining (20) with the observation $\|u_i\|\leq\|u_i\|_M\|M^{-1}\|^{\frac{1}{2}}$.

We next turn to the proof of (19), and for that purpose we remark that $u_i+z\notin u_1\mathbb{Z}+\cdots+u_{i-1}\mathbb{Z}$. Indeed otherwise u_i would be a linear combination of $u_1,\cdots,u_{i-1},u_{i+1},\cdots,u_d$, which contradicts (15). Definition 1.4 thus implies that

$$\|u_i\|_M^2\leq\|u_i+z\|_M^2=\|u_i\|_M^2+2u_i^{\mathrm{T}}Mz+\|z\|_M^2,$$

which implies that $-2u_i^{\mathrm{T}}Mz\leq\|z\|_M^2$. We obtain likewise $2u_i^{\mathrm{T}}Mz\leq\|z\|_M^2$, which concludes the proof of (19). □

We say that point z is a vertex of a mesh $\mathcal{T}$, if it is a vertex of one of the simplices $T \in \mathcal{T}$. We introduce a distance $d_\times$ on the collection S_d^+ of symmetric positive definite matrices, which is defined as follows: for all $M, N \in S_d^+$

$$d_\times(M,N) := \sup_{u \neq 0} |\ln \|u\|_M - \ln \|u\|_N| \,.$$

The distance $d_\times$ allows to compare the norms of vectors multiplicatively, justifying the $\times$ subscript, in contrast with the classical operator norm which is tailored for additive comparisons. Indeed denoting $\alpha := d_\times(M,N)$ and $\beta := \|M^{\frac{1}{2}} - N^{\frac{1}{2}}\|$, one has for all $u \in \mathbb{R}^2$ such that $\|u\| = 1$

$$e^{-\alpha} \leq \|u\|_M / \|u\|_N \leq e^{\alpha}, \quad \text{and} \quad -\beta \leq \|u\|_M - \|u\|_N \leq \beta.$$

The next lemma establishes a lower bound on the $\|\cdot\|_M$ norm of points outside of a N-reduced mesh, when the matrices $M, N \in S_d^+$ are close enough.

Lemma 1.7. *Assume that $1 \leq d \leq 4$. Let $M, N \in S_d^+$. Let $u_1, \cdots, u_d$ be an arbitrary M-reduced basis of $\mathbb{Z}^d$, and let $\mathcal{T}$ be a N-reduced mesh. Consider a point $z \in \mathbb{Z}^d$ which is not a vertex of $\mathcal{T}$. Then there exists $1 \leq l \leq d$ such that*

$$z \in u_1\mathbb{Z} + \cdots + u_l\mathbb{Z} \quad \textit{and} \quad \|z\|_M^2 \, e^{4d_\times(M,N)} \geq \|u_l\|_M^2 + \|u_1\|_M^2.$$

Proof. Since the union of the elements of $\mathcal{T}$ is a neighborhood of the origin, there exists a simplex $T \in \mathcal{T}$ and a real $\lambda > 0$ such that $\lambda z \in T$. Denoting by $v_1, \cdots, v_d$ the non-zero vertices of T, there exists therefore non-negative reals $\alpha_1, \cdots, \alpha_d \in \mathbb{R}_+$ such that $z = \alpha_1 v_1 + \cdots + \alpha_d v_d$. Since $|\det(v_1, \cdots, v_d)| = d!|T| = 1$, the coefficient α_i is an integer for all $1 \leq i \leq d$.

Up to reordering the vertices $v_1, \cdots, v_d$, we may assume that $v_1, \cdots, v_k$ are positive, and $v_{k+1}, \cdots, v_d$ are zero, for some $1 \leq k \leq d$. We denote by l the smallest integer such that $v_1, \cdots, v_k \in u_1\mathbb{Z} + \cdots + u_l\mathbb{Z}$. By additivity, $z \in u_1\mathbb{Z} + \cdots + u_l\mathbb{Z}$. By assumption there exists $1 \leq i \leq k$ such that $v_i \notin u_1\mathbb{Z} + \cdots + u_{l-1}\mathbb{Z}$, and thus $\|v_i\|_M \geq \|u_l\|_M$; other vertices satisfy $\|u_j\|_M \geq \|u_1\|_M$, $1 \leq j \leq k$, since they are non-zero and have integer coordinates. We thus obtain

$$\begin{aligned}
e^{4d_\times(M,N)} \|z\|_M^2 &\geq e^{2d_\times(M,N)} \|z\|_N^2 \\
&= e^{2d_\times(M,N)} \left(\sum_{1 \leq i \leq k} \alpha_i^2 \|v_i\|_N^2 + 2 \sum_{1 \leq i < j \leq k} \alpha_i \alpha_j v_i^{\mathrm{T}} N v_j \right) \\
&\geq e^{2d_\times(M,N)} \sum_{1 \leq i \leq k} \alpha_i^2 \|v_i\|_N^2 \\
&\geq \sum_{1 \leq i \leq k} \alpha_i^2 \|v_i\|_M^2 \\
&\geq \|u_l\|_M^2 + \left(\sum_{2 \leq i \leq k} \alpha_i^2 - 1 \right) \|u_1\|_M^2.
\end{aligned}$$

We have $\alpha_1^2 + \cdots + \alpha_k^2 \geq 2$ since z is not a vertex of T, which concludes the proof. ∎

The following corollary shows that the vertices of a N-reduced mesh contain a M-reduced basis, whenever M and N are close enough. This property will be used in the convergence analysis of the FM-LBR, Lemmas 2.6 and 2.7, to show the connectivity of the graph defined by the local stencils in the discrete domain. It is is illustrated on Figures 2 and 3: the basis vectors u, v are always among the vertices of the consecutive mesh, and conversely.

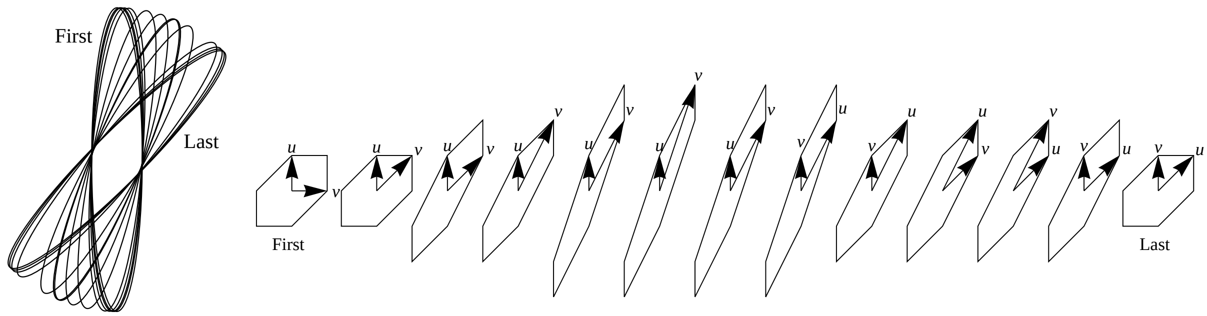

Figure 3: The unit sphere $\{u;\ \|u\|_M = 1\}$, a M-reduced basis (u, v), and the boundary of a M-reduced mesh constructed as in Proposition 1.9, for some $M \in S_2^+$ of anisotropy ratio $\kappa(M) = 6$, and eigenvector $(\cos(\theta), \sin(\theta))$, $\theta \in [\pi/4, \pi/2]$, associated to the small eigenvalue.

Corollary 1.8. *Assume that* $1 \leq d \leq 4$. *Let* $M, N \in S_d^+$ *be such that*

$$d_\times(M, N) < \ln(1 + \kappa(M)^{-2})/4. \tag{21}$$

Let $(u_1, \cdots, u_d)$ *be a* M*-reduced basis of* $\mathbb{Z}^d$, *and let* $\mathcal{T}$ *be a* N*-reduced mesh. Then* $u_1, \cdots, u_d$ *and* $-u_1, \cdots, -u_d$ *are vertices of* $\mathcal{T}$.

Proof. We consider $1 \leq l \leq d$, and we assume for contradiction that u_l (or $-u_l$) is not a vertex of $\mathcal{T}$. It follows from (21) and (18) that

$$\|u_l\|_M^2 e^{4d_\times(M,N)} < \|u_l\|_M^2 + \kappa(M)^{-2}\|u_l\|_M^2 \leq \|u_l\|_M^2 + \|u_1\|_M^2.$$

The previous lemma thus implies $u_l \in u_1\mathbb{Z} + \cdots + u_k\mathbb{Z}$ for some $1 \leq k < l$. This contradicts Definition 1.4, and concludes the proof. $\square$

1.2 Explicit construction of M-reduced meshes

This subsection is devoted to the explicit construction of a M-reduced mesh of bounded cardinality for any $M \in S_d^+$, where $d = 2$ in Proposition 1.9, and $d = 3$ in Proposition 1.10. This construction uses as a starting point a M-reduced basis of $\mathbb{Z}^d$.

Let A be a $d \times d$ invertible matrix. For any simplex T we denote $A(T) := \{Az;\ z \in T\}$, and for any mesh $\mathcal{T}$ we denote $A(\mathcal{T}) := \{A(T);\ T \in \mathcal{T}\}$.

Proposition 1.9. *Let* $M \in S_2^+$ *and let* (u_1, u_2) *be a* M*-reduced basis of* $\mathbb{Z}^2$. *Let* $u := u_1$ *and let* $v := \varepsilon u_2$, *where* $\varepsilon \in \{-1, 1\}$ *is chosen in such way that* $u^{\mathrm{T}} M v \geq 0$. *Consider the collection* $\mathcal{T}$ *of triangles of vertices* 0 *and one of the following pairs*

$$\{u, v\}, \{v, v - u\}, \{v - u, -u\}$$

or their opposites (the opposite of the pair $\{a, b\}$ *is* $\{-a, -b\}$*). Then* $\mathcal{T}$ *is a* M*-reduced mesh.*

Proof. We denote by (a),(b),(c) the corresponding points of the Definition 1.1 of M-reduced meshes.

Consider a matrix A such that $Au = (1, 0)$ and $Av = (0, 1)$. Then the mesh $A(\mathcal{T})$ does not depend on u and v, and the reunion of its elements is easily checked to be a neighborhood of the origin. This immediately implies (a).

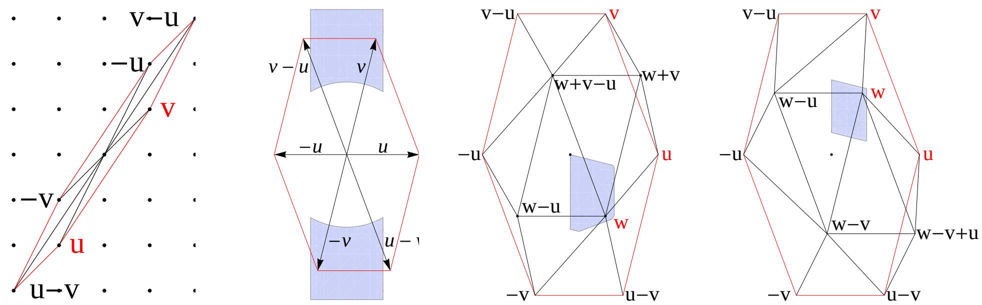

Figure 4: A M-reduced mesh for particular a matrix $M \in S_2^+$, such that $\kappa(M) = 10$ (left). Illustration for the proof of Proposition 1.9 (center left). Connectivity defined by (23) (center right) and by (25) (right), in Proposition 1.10.

The equality of the determinants

$$\det(u, v) = \det(v, v - u) = \det(v - u, -u)$$

shows that all the elements of $\mathcal{T}$ have volume 1/2. Since their vertices clearly lie on $\mathbb{Z}^2$, this establishes (b).

We have by construction $u^{\mathrm{T}} M v \geq 0$. Furthermore, applying (19) to the vectors u_1 and $z = u_2$ (resp. u_2 and $z = u_1$) we obtain $2|u^{\mathrm{T}} M v| \leq \|v\|_M^2$ and $2|u^{\mathrm{T}} M v| \leq \|u\|_M^2$, hence

$$v^{\mathrm{T}} M(v - u) = \|v\|_M^2 - u^{\mathrm{T}} M v \geq 0, \qquad (v - u)^{\mathrm{T}} M(-u) = \|u\|_M^2 - u^{\mathrm{T}} M v \geq 0,$$

which establishes point (c) and concludes the proof.

For a visual illustration of point (c) we suggest the reader to look at Figure 4 (center left), which displays the *image* of the mesh $\mathcal{T}$ by a linear transformation P such that $Pu = (1, 0)$ and $P^{\mathrm{T}} P = M$. For any two vectors $a, b \in \mathbb{R}^2$, we have $a^{\mathrm{T}} M b \geq 0$ if and only if the images of a and b by P form an acute angle. The blue region corresponds to all possible values of v which satisfy the constraints $2|v^{\mathrm{T}} M u| \leq \|u\|_M^2 \leq \|v\|_M^2$. □

Proposition 1.10. *Let $M \in S_3^+$ and let (u_1, u_2, u_3) be a M-reduced basis of $\mathbb{Z}^3$. We distinguish two cases depending on the parity of the number of non-negative scalar products among $u_1^{\mathrm{T}} M u_2$, $u_2^{\mathrm{T}} M u_3$, $u_3^{\mathrm{T}} M u_1$.*

- *Odd parity. We denote $u := \varepsilon_1 u_1$, $v := \varepsilon_2 u_2$, $w := \varepsilon_3 u_3$, where $\varepsilon_1, \varepsilon_2, \varepsilon_3 \in \{-1, 1\}$ are chosen in such way that*

$$u^{\mathrm{T}} M v \geq 0, \quad u^{\mathrm{T}} M w \geq 0, \quad v^{\mathrm{T}} M w \leq 0. \tag{22}$$

Consider the collection $\mathcal{T}$ of tetrahedra of vertices 0 and one of the following triplets

$$\begin{array}{ccc} \{w, u, w + v\} & \{w, w + v, w + v - u\} & \{w, w + v - u, w - u\} \\ \{w, w - u, -v\} & \{w, -v, u - v\} & \{w, u - v, u\} \\ \{v, w + v, u\} & \{v, w + v - u, w + v\} & \{v, v - u, w + v - u\} \\ \{-u, -v, w - u\} & \{-u, w - u, w + v - u\} & \{-u, w + v - u, v - u\} \end{array} \tag{23}$$

or their opposites (the opposite of the triplet $\{a, b, c\}$ is $\{-a, -b, -c\}$). Then $\mathcal{T}$ is a M-reduced mesh.

- *Even parity. We denote $u := \varepsilon_1 u_{\sigma_1}$, $u := \varepsilon_2 u_{\sigma_2}$, $w := \varepsilon_3 u_{\sigma_3}$, where $\varepsilon_1, \varepsilon_2, \varepsilon_3 \in \{-1, 1\}$ and the permutation σ of $\{1, 2, 3\}$ are chosen in such way that*

$$v^{\mathrm{T}} M w \geq u^{\mathrm{T}} M v \geq u^{\mathrm{T}} M w \geq 0. \tag{24}$$

Consider the collection $\mathcal{T}$ of tetrahedra of vertices 0 and one of the following triplets

$$\begin{array}{ccc} \{u, v, w\} & \{w, v, w-u\} & \{w-u, v, v-u\} \\ \{w-u, v-u, -u\} & \{w-u, -u, w-v\} & \{w-v, -u, -v\} \\ \{w-v, -v, u-v\} & \{w-v, u-v, w-v+u\} & \{w-v+u, u-v, u\} \\ \{u, w, w-v+u\} & \{w, w-v, w-v+u\} & \{w, w-u, w-v\} \end{array} \tag{25}$$

or their opposites (the opposite of the triplet $\{a, b, c\}$ is $\{-a, -b, -c\}$. Then $\mathcal{T}$ is a M-reduced mesh.

Proof. We denote by (a),(b),(c) the corresponding points of the Definition 1.1 of M-reduced meshes.

In each case of the parity, if A is a matrix such that $Au = (1, 0, 0)$, $Av = (0, 1, 0)$ and $Aw = (0, 0, 1)$, then the mesh $A(\mathcal{T})$ is independent of u, v, w. The union of the elements of this constant mesh $A(\mathcal{T})$ is easily seen to be a neighborhood of the origin, see Figure 4, which implies (a).

The determinant of each triplet of vectors appearing in (23) and (25), is equal to $\det(u, v, w)$. Since $|\det(u, v, w)| = 1$, see (15), the volume of all the constructed simplices is $1/6$. Furthermore their vertices obviously lie on $\mathbb{Z}^3$, which establishes (b).

We next need to establish (c): the non-negativeness of the scalar products between two common vertices of any tetrahedron $T \in \mathcal{T}$. To avoid notational clutter we denote in this proof

$$\langle u, v \rangle := u^{\mathrm{T}} M v.$$

We first remark that the positivity of the scalar products $\langle u, v\rangle$, $\langle v, v-u\rangle$,$\langle v-u, -u\rangle$ (associated to edges which lie in the plane generated by (u, v)) follows as in Proposition 1.9 from the inequalities $\langle u, v\rangle \geq 0$, $2|\langle u, v\rangle| \leq \|u\|_M^2$ and $2|\langle u, v\rangle| \leq \|v\|_M^2$. For the other scalar products we need to distinguish between the two cases of the parity.

Odd parity. It follows from (22) that

$$\langle u, v\rangle \geq 0,\ \langle u, w\rangle \geq 0,\ \langle -v, w\rangle \geq 0,$$
$$\langle u, w+v\rangle \geq 0,\ \langle w, u-v\rangle \geq 0,\ \langle w-u, -v\rangle \geq 0.$$

Applying (19) to $z = u$ we obtain $2|\langle u, w\rangle| \leq \|u\|_M^2$, hence $\langle -u, w-u\rangle = \|u\|_M^2 - \langle u, w\rangle \geq 0$. Likewise

$$\langle v, w+v\rangle \geq 0, \langle w, w+v\rangle \geq 0, \langle w, w-u\rangle \geq 0$$

Applying (19) to $z = u$ we obtain $2|\langle u, v\rangle| \leq \|u\|_M^2$ and $2|\langle u, w\rangle| \leq \|u\|_M^2$, which implies $\langle -u, w+v-u\rangle = \|u\|_M^2 - \langle u, v\rangle - \langle u, w\rangle \geq 0$. Likewise

$$\langle v, w+v-u\rangle \geq 0,\ \langle w, w+v-u\rangle \geq 0.$$

Applying (19) to $z = v-u$ we obtain $2|\langle w, v-u\rangle| \leq \|v-u\|_M^2$, hence $\langle v-u, w+v-u\rangle = \|v-u\|_M^2 + \langle w, v-u\rangle \geq 0$. Likewise

$$\langle w-u, w+v-u\rangle \geq 0,\ \langle w+v, w+v-u\rangle \geq 0.$$

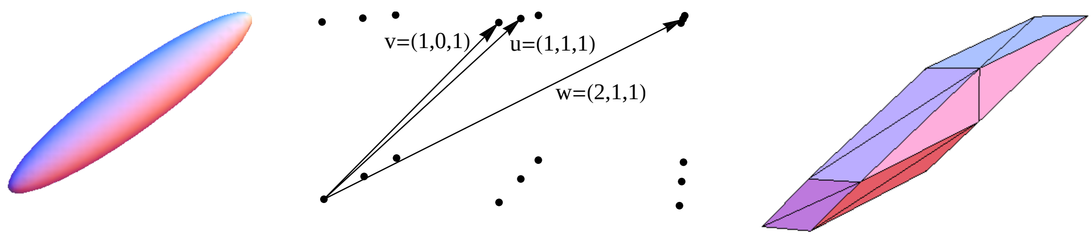

Figure 5: The unit sphere $\{u;\ \|u\|_M = 1\}$, a M-reduced basis (u, v, w), and a M-reduced mesh constructed as in Proposition 1.10, for the symmetric matrix $M \in S_3^+$ of eigenvalues $5^2, 5^2, 1$ (anisotropy ratio $\kappa(M) = 5$) and eigenvector $(4, 2, 3)$ associated to the small eigenvalue.

Even parity. From the inequality (24) we obtain

$$\langle v, w\rangle \geq \langle u, v\rangle \geq \langle u, w\rangle \geq 0,\ \ \langle v, w-u\rangle \geq 0,\ \ \langle -u, w-v\rangle \geq 0.$$

Applying (19) to $z = u$ we obtain $2|\langle u, v\rangle| \leq \|u\|_M^2$, and therefore $\langle -u, w-u\rangle \geq 0$. Likewise

$$\langle -v, w-v\rangle \geq 0,\ \langle w, w-u\rangle \geq 0,\ \langle w, w-v\rangle \geq 0.$$
$$\langle u, w-v+u\rangle \geq 0,\ \langle w, w-v+u\rangle \geq 0.$$

Applying (19) to $z = u$ we obtain $2|\langle u, v\rangle| \leq \|u\|_M^2$ and $2|\langle u, w\rangle| \leq \|u\|_M^2$, which implies $\langle v-u, w-u\rangle = \langle v, w\rangle + (\|u_M\|^2 - \langle u, v\rangle - \langle u, w\rangle) \geq 0$. Likewise

$$\langle w-u, w-v\rangle \geq 0,\ \langle w-v, u-v\rangle \geq 0.$$

Applying (19) to $z = u - v$ we obtain $|\langle w, u-v\rangle| \leq \|u-v\|_M^2$, hence $\langle u-v, w-v+u\rangle \geq 0$. Likewise $\langle w-v, w-v+u\rangle \geq 0$. $\square$

2 The algorithm

We present in this section our modified fast marching algorithm, Fast Marching using Lattice Basis Reduction (FM-LBR). We estimate its complexity, and we establish its consistency: the convergence of the discrete approximations towards the solution of the continuous problem. Our presentation of this section is fairly classical, and we invoke (variants of) arguments previously seen in the literature on numerical schemes for the eikonal equation [22, 19, 5]. A detailed description and proof is nevertheless necessary, since the original constructions of §1 do not exactly fit in earlier framework. The FM-LBR is split in two steps: Preprocessing (stencil construction), followed by Execution (Dijkstra-like resolution of (10)). The integer $n \geq 2$ is fixed in this subsection.

After the Preprocessing of the FM-LBR comes its proper Execution. This part of the algorithm involves a boolean table $b : \Omega_n \to \{\mathit{trial}, \mathit{accepted}\}$. Given such a table, we say that a grid point $z \in \Omega$ is accepted if $b(z) = \mathit{accepted}$, and that z is a trial point otherwise.

Definition 2.1. *Consider a map* $\mathrm{d} : \Omega_n \to \overline{\mathbb{R}}_+$, *and a grid point* $x \in \Omega_n$. *We introduce a variant of the Hopf-Lax update* $\Lambda_n(\mathrm{d}, x)$ *which involves two additional variables: a boolean table* $b : \Omega_n \to \{\mathit{trial}, \mathit{accepted}\}$, *and another grid point* $y \in \Omega_n$.

$$\Lambda_n(\mathrm{d}, x;\, b, y) := \min_{x' \in \Gamma} \|x' - x\|_{\mathcal{M}(x)} + \mathrm{I}_V\, d(x'). \tag{26}$$

FM-LBR: Preprocessing.

Input. An integer $n \geq 2$. The values $\mathcal{M}(z)$, $z \in \Omega_n$, of a Riemannian metric $\mathcal{M} \in C^0(\Omega, S_d^+)$.

1. Produce a $\mathcal{M}(x)$-reduced mesh $\mathcal{T}(x)$ for each $x \in \Omega_n$, using Proposition 1.9 in dimension 2, Proposition 1.10 in dimension 3, or by combining these constructions with Remark 1.3 in higher dimension if the metric has a block diagonal structure.
2. Construct the stencils $V_n(x) = x + \frac{1}{n}\mathcal{T}(x)$, $x \in \Omega_n$, by offsetting and rescaling $\mathcal{T}(x)$.
3. Compute the *reversed* stencils, defined by $V^n(y) := \{x \in \Omega_n \setminus \{y\};\ y \text{ is a vertex of } V_n(x)\}$.

We denoted by Γ the collection of vertices, edges and facets of the boundary of the stencil $V := V_n(x)$, which (i) contain only accepted *grid points, and (ii) contain the point y. As before I_V denotes piecewise linear interpolation on the mesh V. The difference $x' - x$ makes sense because x and x' belong to the small and simply connected stencil V of x.*

Equivalently, one has

$$\Lambda_n(\mathrm{d}, x;\, b, y) := \min_{k,(\alpha_i),(v_i)} \left\{ \left\| \sum_{1 \leq i \leq k} \alpha_i \frac{v_i}{n} \right\|_{\mathcal{M}(z)} + \sum_{1 \leq i \leq k} \alpha_i \ \mathrm{d}\left(x + \frac{v_i}{n}\right) \right\}, \tag{27}$$

where the minimum is taken among all $1 \leq k \leq d$, all $\alpha_1, \cdots, \alpha_k \in \mathbb{R}$ and all $v_1, \cdots, v_k \in \mathbb{Z}^d$, such that

- *$\alpha_i + \cdots + \alpha_k = 1$, and $\alpha_i \geq 0$ for all $1 \leq i \leq k$.*
- *$v_1, \cdots, v_k$ are non-zero vertices of a common simplex $T \in \mathcal{T}(x)$.*
- *$b(x + v_1/n) = \cdots = b(x + v_k/n) = accepted$.*
- *$x + v_1/n = y$.*

The main originality of the FM-LBR lies in the stencil construction, in the Preprocessing step. The second part of this algorithm, Execution, is similar to other Dijkstra inspired solvers of the eikonal equation with static stencils, such as [2, 22].

FM-LBR: Execution.

4. Initialize to *trial* a (mutable) boolean table $b : \Omega_n \to \{trial, accepted\}$.
5. Initialize to $+\infty$ another (mutable) table $\mathrm{d} : \Omega_n \to \overline{\mathbb{R}}_+$, except for $\mathrm{d}(0) = 0$.
6. **While** there remains a *trial* point in Ω_n
   (a) Denote by y the *trial* point which minimizes d, and set $b(y) ::= accepted$.
   (b) For all $x \in V^n(y)$, set $\mathrm{d}(x) ::= \min\{\mathrm{d}(x),\, \Lambda_n(\mathrm{d}, x;\, b, y)\}$.

Output. The map $\mathrm{d} : \Omega_n \to \overline{\mathbb{R}}_+$.

The rest this section is devoted to the proof of the following theorem. Points (i) and (iii) also hold for any alternative construction of $\mathcal{M}(z)$-reduced meshes $\mathcal{T}(z)$, in the Preprocessing step, in dimension $d \leq 4$, which diameters are uniformly bounded by a constant $C = C(\kappa(\mathcal{M}))$.

Theorem 2.2. *(i) The map* $\mathrm{d}_n = \mathrm{d}$ *obtained at the end of the FM-LBR satisfies the discrete fixed point system (10).*

(ii) The cost of the FM-LBR is $\mathcal{O}(N \ln N + N \ln \kappa(\mathcal{M}))$*, where* $N := n^d = \#(\Omega_n)$*, if elementary operations among reals* $(+, -, \times, /$ *and* $\sqrt{\cdot})$ *have unit complexity.*

(iii) Denoting by $\mathrm{D} : \Omega \to \mathbb{R}_+$ *the solution of the anisotropic eikonal equation (2) one has*

$$\lim_{n\to\infty} \left(\max_{z \in \Omega_n} |\, \mathrm{D}(z) - \mathrm{d}_n(z)| \right) = 0.$$

The next lemma constitutes the causality property, which is the key to the one pass resolution of the system (10) in Fast Marching algorithms, see Proposition 2.4.

Lemma 2.3 (Sethian and Vladimirsky 2000, [19]). *Let* $M \in S_d^+$ *and let* $v_1, \cdots, v_k \in \mathbb{R}^d$ *be linearly independent vectors such that* $v_i^{\mathrm{T}} M v_j \geq 0$ *for all* $1 \leq i \leq j \leq k$*. Consider the compact and convex set*

$$\Xi := \{\alpha = (\alpha_1, \cdots, \alpha_k) \in \mathbb{R}_+^k;\ \alpha_1 + \cdots + \alpha_k = 1\}$$

Choose $\Delta = (\delta_1, \cdots, \delta_k) \in \mathbb{R}$*, denote*

$$\delta := \min_{\alpha \in \Xi} \left\| \sum_{1 \leq i \leq k} \alpha_i w_i \right\|_M + \sum_{1 \leq i \leq k} \alpha_i \delta_i, \tag{28}$$

and assume that the minimizer $\alpha \in \Xi$ *of this problem has positive coefficients.*

Then $\delta > \delta_i$ *for each* $1 \leq i \leq k$*. Furthermore denoting by* $\hat{M}$ *the matrix of entries* $\hat{M}_{ij} = v_i^{\mathrm{T}} M v_j$*, and defining* $\mathbf{1} := (1, \cdots, 1) \in \mathbb{R}^k$*, one has the relations*

$$1 = \|\delta \mathbf{1} - \Delta\|_{\hat{M}^{-1}}, \tag{29}$$

$$\hat{M}\alpha = \|\alpha\|_{\hat{M}} (\delta \mathbf{1} - \Delta). \tag{30}$$

Proof. For completeness, and due to notation changes, we give the proof of this lemma, which follows the steps of Property A.1 in [19].

The problem (28) is the optimization, on the compact an convex set Ξ, of a strictly convex functional. Hence there exists a unique minimizer $\alpha \in \Xi$. The gradients of the maps $\alpha \mapsto \|\alpha\|_{\hat{M}} + \alpha^{\mathrm{T}}\Delta$ and $\alpha \mapsto \alpha^{\mathrm{T}} 1$, respectively the minimized function and the constraint in (28), are proportional at the minimizer α according to Lagrange's Theorem. Hence there exists $\lambda \in \mathbb{R}$, the Lagrange multiplier, such that

$$\frac{\hat{M}\alpha}{\|\alpha\|_{\hat{M}}} + \Delta = \lambda \mathbf{1}. \tag{31}$$

Therefore

$$\|\lambda \mathbf{1} - \Delta\|_{\hat{M}^{-1}} = \left\| \frac{\hat{M}\alpha}{\|\alpha\|_{\hat{M}}} \right\|_{\hat{M}^{-1}} = 1. \tag{32}$$

Using again (31) we obtain

$$\delta = \|\alpha\|_{\hat{M}} + \alpha^{\mathrm{T}}\Delta = \alpha^{\mathrm{T}} \left(\frac{\hat{M}\alpha}{\|\alpha\|_{\hat{M}}} + \Delta \right) = \alpha^{\mathrm{T}}(\lambda \mathbf{1}) = \lambda.$$

Combining the last equation with (32) (resp. (31)) we obtain (29) (resp. (30)).

Since $\hat{M}$ has non negative coefficients, positive diagonal coefficients, and since we have assumed that α has positive coefficients, we obtain that the product $\hat{M}\alpha$ has positive coefficients, hence also $\delta\mathbf{1} - \Delta$. Therefore $\delta > \delta_i$ for all $1 \leq i \leq k$ in view of (30), which concludes the proof. $\square$

The next proposition, adapted from [22], shows that the Execution of the FM-LBR solves the discrete fixed point system (10).

Proposition 2.4 (Tsitsiklis, 1995)**.** *We denote by E the collection of maps $\mathrm{d} : \Omega_n \to \overline{\mathbb{R}}_+$ satisfying $\mathrm{d}(0) = 0$. The set E is partially ordered as follows: $\mathrm{d} \leq \mathrm{d}'$ if $\mathrm{d}(z) \leq \mathrm{d}'(z)$ for all $z \in \Omega_n$. We denote by $\mathrm{d}^0 \in E$ the map equal to $+\infty$ on $\Omega_n \setminus \{0\}$.*

We denote by $\Lambda_n : E \to E$ the operator defined by $\Lambda_n(\mathrm{d})(z) := \Lambda_n(\mathrm{d}, z)$ if $z \in \Omega_n \setminus \{0\}$.

1. *The operator Λ_n is increasing on E: $\Lambda_n(\mathrm{d}) \leq \Lambda_n(\mathrm{d}')$ for all $\mathrm{d}, \mathrm{d}' \in E$ such that $\mathrm{d} \leq \mathrm{d}'$.*

2. *The sequence $\Lambda_n^k(\mathrm{d}^0)$ converges to a fixed point $\overline{\mathrm{d}} \in E$ of Λ_n, as $k \to \infty$.*

Let $N := n^d$. We denote by $(y_i)_{0 \leq i < N}$ the grid points consecutively selected at step 6 (a) of the FM-LBR, and by $(\mathrm{d}^i)_{1 \leq i \leq N}$ the maps obtained at the end of each iteration of the while loop.

3. *For all $0 \leq i \leq j \leq N$ one has $\mathrm{d}^i \geq \mathrm{d}^j \geq \overline{\mathrm{d}}$.*

4. *For all $0 \leq i < N$ one has $\mathrm{d}^j(y_i) = \overline{\mathrm{d}}(y_i)$.*

5. *For all $0 \leq i \leq j < N$ one has $\overline{\mathrm{d}}(y_i) \leq \overline{\mathrm{d}}(y_j)$.*

Combining Points 3 and 4 we obtain $\mathrm{d}^N = \overline{\mathrm{d}}$. Recalling Point 2 and observing that a fixed point of Λ_n solves (10), we obtain Point (i) of Theorem 2.2.

Proof. Point 1, the monotonicity of Λ_n on E, follows from the expression (9) of the pointwise entries $\Lambda_n(\mathrm{d}, z)$, $\mathrm{d} \in E$, $z \in \Omega_n$, as a minimum of monotone maps (linear forms with positive coefficients).

Point 2. Any map $\mathrm{d} \in E$ satisfies $\mathrm{d} \leq \mathrm{d}^0$, hence $\Lambda(\mathrm{d}^0) \leq \mathrm{d}^0$. Using Point 1 and an immediate induction argument we obtain that the sequence $\Lambda^k(\mathrm{d}^0)$ is decreasing, and thus converging to a limit $\overline{\mathrm{d}} \in E$. The identity $\Lambda(\overline{\mathrm{d}}) = \overline{\mathrm{d}}$ follows from the continuity of Λ_n, which is clear in view of its expression (9) as a minimum of linear functions indexed by a compact set.

Point 3. Consider a fixed $0 \leq i < N$, and denote by b_i the value of the boolean map at the beginning of the i-th iteration (thus $b_i(y_j) = \mathit{trial}$ if and only if $i \leq j$). The maps d^i and d^{i+1} may only differ on the set $V^n(y_i)$, see step 6 (b) of the FM-LBR. We assume that $\mathrm{d}^i \geq \overline{\mathrm{d}}$, an induction hypothesis that is satisfied for $i = 0$, and observe that $\mathrm{d}^i \geq \mathrm{d}^{i+1}$ by construction. Furthermore, for $x \in V^n(y_i)$ we have either $\mathrm{d}^{i+1}(x) = \mathrm{d}^i(x)$, hence $\mathrm{d}^{i+1}(x) \geq \overline{\mathrm{d}}(x)$ by induction, or

$$\mathrm{d}^{i+1}(x) = \Lambda_n(\mathrm{d}^i, x; b_i, y_i) \geq \Lambda_n(\overline{\mathrm{d}}, x; b_i, y_i) \geq \Lambda_n(\overline{\mathrm{d}}, x) = \overline{\mathrm{d}}(x).$$

Thus $\mathrm{d}^i \geq \mathrm{d}^{i+1} \geq \overline{\mathrm{d}}$, for all $0 \leq i < N$, which concludes the proof of Point 3.

Following the steps of Lemma 3.2 in [22], we prove Points 4 and 5, simultaneously, by induction on i. Let y be the trial point ($b_i(y) = \mathit{trial}$, equivalently $y = y_j$ for some $j \geq i$) which minimizes $\overline{\mathrm{d}}$. If $y \in \partial\Omega_* = \{0\}$, then $y = 0$, $i = 0$, $\overline{\mathrm{d}}(0) = \mathrm{d}^0(0) = 0$, and there is nothing to prove. Otherwise we have

$$\overline{\mathrm{d}}(y) = \Lambda_n(\overline{\mathrm{d}}, y) = \|y - \sum_{1 \leq r \leq k} \alpha_r y_{\sigma(r)}\|_{\mathcal{M}(y)} + \sum_{1 \leq r \leq k} \alpha_r \overline{\mathrm{d}}(y_{\sigma(r)}), \tag{33}$$

for some $1 \leq k \leq d$, positive barycentric coefficients $(\alpha_r)_{r=1}^k$, and grid points $(y_{\sigma(r)})_{r=1}^k$. These points satisfy $\overline{\mathrm{d}}(y_{\sigma(r)}) < \overline{\mathrm{d}}(y)$, due to the causality property, hence $b_i(y_{\sigma(r)}) = \mathit{accepted}$. Thus $\sigma(r) < i$, for $1 \leq r \leq k$, and therefore $\mathrm{d}^{i-1}(y_{\sigma(r)}) = \overline{\mathrm{d}}(y_{\sigma(r)})$ by induction hypothesis on Point 4. Recalling the Hopf-Lax updates performed at step 6 (b) of the FM-LBR, in the previous iterations of the while loop, we find that the right hand side of (33) is an upper bound for $\mathrm{d}^i(y)$, hence $\mathrm{d}^i(y) \leq \overline{\mathrm{d}}(y)$. Since $\mathrm{d}^i \geq \mathrm{d}$ (Point 3), it follows that y minimizes d^i among trial points, similarly to y_i. Finally $\overline{\mathrm{d}}(y) \geq \mathrm{d}^i(y) = \mathrm{d}^i(y_i) \geq \overline{\mathrm{d}}(y_i) \geq \overline{\mathrm{d}}(y)$, using Point 3 and the optimality of y for the last two inequalities. Thus $\overline{d}(y_i) = \mathrm{d}^i(y_i)$, and this is the minimum of $\overline{\mathrm{d}}$ among the trial points $\{y_j;\, j \geq i\}$. This establishes Points 4 and 5 at rank i. $\square$

We next turn to Point (ii) of Theorem 2.2: the complexity estimate of the FM-LBR. Consider a point $z \in \Omega_n$. Obtaining a $\mathcal{M}(z)$-reduced basis has cost $\mathcal{O}(1 + \ln \kappa(\mathcal{M}(z)))$, see Theorem 1.5, and from this point constructing a $\mathcal{M}(z)$-reduced mesh $\mathcal{T}(z)$ using Proposition 1.9 or 1.10 has cost $\mathcal{O}(1)$. Step 1 of FM-LBR Preprocessing thus has total cost $\mathcal{O}(N + N \ln \kappa(z))$. Steps 2 and 3 only cost $\mathcal{O}(N)$.

The FM-LBR Execution involves the evaluation of the Hopf-Lax update operator (27). Each mesh $\mathcal{T}(z)$, $z \in \Omega_n$, has $\mathcal{O}(1)$ vertices by assumption, hence evaluating this operator amounts to solve $\mathcal{O}(1)$ optimization problems of the form (28). Their solution is obtained as the root δ of a univariate quadratic equation (29), hence the complexity is $\mathcal{O}(1)$. Globally in the Execution, such evaluations happen at most the following number of times:

$$\sum_{z \in \Omega_n} \#(V^n(z)) = \sum_{z \in \Omega_n} n(z) = \mathcal{O}(N),$$

where $n(z)$ denotes the number of non-zero vertices of $\mathcal{T}(z)$. In our case $n(z) = \mathcal{O}(1)$. In addition, the FM-LBR Execution requires to maintain a list of the elements of $\beta^{-1}(\mathit{trial})$, sorted by increasing value of δ. Elementary insertions and deletions in this list have cost $\mathcal{O}(\ln N)$ (if this sorted list is implemented numerically using a binary heap), and occur each time the Hopf-Lax formula is evaluated. The Execution cost is thus $\mathcal{O}(N \ln N)$. Combining the costs of Preprocessing and Execution, we obtain as announced the complexity $\mathcal{O}(N \ln N + N \ln \kappa(\mathcal{M}))$.

Remark 2.5 (Memory requirements)**.** *The memory requirements of numerical methods for the eikonal equation, such as the AGSI and the FM-LBR, are dominated by (I) storing the riemannian metric $\mathcal{M}$, sampled on the discrete domain Ω_*, and the discrete solution* d*, and (II) storing the graph structure underlying the numerical scheme. Point (I) requires two tables of $Nd(d+1)/2$ and N reals, typically represented in 64bit floating point format, independently of the method.*

Point (II) is avoided for AGSI when this method is executed on a mesh with a trivial periodic structure, which is the case in our experiments. For the FM-LBR, Point (II) is divided into the storage of the direct stencils $(V_*(x))_{x \in \Omega_n}$, and of the reversed stencils $(V^*(y))_{y \in \Omega_n}$. Storing a directed graph with N vertices and N' edges requires a table of N' vectors[1]*, and another of N integers: the indices of the start of the consecutive stencils in the previous table. For the direct or the reversed stencils of the FM-LBR, $N' = 6N$ in two dimensions, and $N' = 14N$ in three dimensions, see Propositions 1.9 and 1.10. We represented integers in 32bit format, and vector components in 8bit format, since these are small integers by construction. Summing up, we find that the memory requirements of the FM-LBR are approximately twice those of the AGSI (for which we do not count any mesh structure).*

[1] The edge $\vec{xx'}$ may also be represented by the index of the endpoint x', instead of the associated vector.

2.1 Convergence analysis

We establish in this section the convergence result announced in Point (iii) of Theorem 2.2. The FM-LBR is applied for each $n \geq 2$ to a fixed Riemannian metric $\mathcal{M} \in C^0(\Omega, S_d^+)$, and produces a sequence of maps $\mathrm{d}_n : \Omega_n \to \overline{\mathbb{R}}_+$ satisfying (10).

Lemma 2.6. • *(Limited extension) There exists a constant $V_0 = V_0(\mathcal{M})$ such that each vertex v of each mesh $\mathcal{T}(z)$, $z \in \Omega$, satisfies $\|v\| \leq V_0$.*

- *(Consistency) There exists a constant $r_0 = r_0(\mathcal{M}) > 0$ such that the following holds. For each $z \in \Omega$ there exists a basis $(u_1, \cdots, u_d)$ of $\mathbb{Z}^d$ such that $u_1, \cdots, u_d, -u_1, \cdots, -u_d$ are vertices of the mesh $\mathcal{T}(z + w)$, for all $w \in \mathbb{R}^d$ such that $\|w\| \leq r_0$.*

Proof. We prove this lemma for the mesh constructions given in Propositions 1.9 and 1.10. The "block diagonal" case immediately follows by considering each block separately.

The elements of a M-reduced basis of $\mathbb{Z}^d$ have a norm bounded by $\kappa(M)$, see Proposition 1.6. Hence any vertex of a M-reduced mesh $\mathcal{T}$ built as described in Proposition 1.9 satisfies $\|v\| \leq 2\kappa(M)$ (resp. Proposition 1.10 and $\|v\| \leq 3\kappa(M)$). This establishes Point (Consistency).

We denote by $r_0 = r_0(\mathcal{M})$ the largest positive constant such that for all $z \in \Omega$ and all $w \in \mathbb{R}^d$

$$|w| \leq r_0 \ \Rightarrow \ d_\times(\mathcal{M}(z), \mathcal{M}(z + w)) \leq \ln(1 + \kappa(\mathcal{M})^{-2})/4.$$

If $z \in \Omega$ and $\|w\| \leq r_0$, then it follows from Corollary 1.8 that the vertices of $\mathcal{T}(z + w)$, which is $\mathcal{M}(z + w)$-reduced mesh, contain as a subset any $\mathcal{M}(z)$-reduced basis $(u_1, \cdots, u_d)$ of $\mathbb{Z}^d$. This establishes Point (Limited extension), and concludes the proof. $\square$

Our next lemma shows that the discrete maps $\mathrm{d}_n : \Omega_n \to \overline{\mathbb{R}}_+$ obey a Lipschitz regularity property, if n is sufficiently large, and thus take finite values,

Lemma 2.7. *For any $n > 0$, any $z \in \Omega_n$ and any vertex v of $\mathcal{T}(z)$, one has*

$$\mathrm{d}_n(z) - \mathrm{d}_n(z + v/n) \leq \|v\|_{\mathcal{M}(z)}/n. \tag{34}$$

For any $n \geq n_0$, any $z \in \Omega_n$, and any $v \in \{-1, 0, 1\}^d$ one has

$$|\,\mathrm{d}_n(z) - \mathrm{d}_n(z + v/n)| \leq \Delta_0/n, \tag{35}$$

where $n_0 := d^2 V_0^d / r_0$, $\Delta_0 := d^2 V_0^d M_0$, and $M_0 := \max\{\sqrt{\|\mathcal{M}(z)\|};\, z \in \Omega\}$.

Proof. If $z \in \Omega_n \setminus \{0\}$, then (34) follows from the equality $\mathrm{d}_n(z) = \Lambda_n(\mathrm{d}_n, z)$, and from the definition (27) of $\Lambda_n(\mathrm{d}_n, \cdot)$. If $z = 0$, then the left hand side of (34) is negative, while the right hand side is non-negative. This concludes the proof of the first part of this lemma.

We next turn to the proof of (35), and for that purpose we consider a fixed $n \geq n_0$. Consider a point $z \in \Omega_n$ and a basis $(u_1, \cdots, u_d)$ of $\mathbb{Z}^d$ as in Lemma 2.6 (Consistency), and observe that $\|u_i\| \leq V_0$ for all $1 \leq i \leq d$. Let A be the $d \times d$ matrix which columns are $u_1, \cdots, u_d$. We denote by $\mathrm{com}(A)$ the comatrix of A, and we observe that the coefficients of this matrix are bounded in absolute value by V_0^{d-1} (indeed they are determinants of $(d-1) \times (d-1)$ sub-matrices of A, the norm of which columns is bounded by V_0). Since $|\det A| = 1$, the absolute value of the coefficients of $A^{-1} = \mathrm{com}(A)^{\mathrm{T}} / \det A$ is also bounded by V_0^{d-1}.

Let $(\alpha_1, \cdots, \alpha_d) \in \mathbb{Z}^d$ be such that $v = \alpha_1 u_1 + \cdots + \alpha_d u_1$, in other words $(\alpha_1, \cdots, \alpha_d) = A^{-1} v$. We denote $s := |\alpha_1| + \cdots + |\alpha_d|$, and we observe that $s \leq d^2 V_0^{d-1}$, since the absolute value of the coefficients of v is bounded by 1. There exists $v_0, v_1, \cdots, v_s \in \mathbb{Z}^d$, such that

$v_0 = 0$, $v_s = v$ and $v_{i+1} - v_i \in \{u_1, \cdots, u_d, -u_1, \cdots, -u_d\}$ for all $1 \leq i \leq s$. Furthermore $\|v_i\| \leq sV_0 \leq d^2 V_0^d$, hence $u_1, \cdots, u_d, -u_1, \cdots, -u_d$ are vertices of $\mathcal{T}(z + v_i/n)$ according to Lemma 2.6 (Consistency). It thus follows from (34) that

$$n(\mathrm{d}_n(z + v_i/n) - \mathrm{d}_n(z + v_{i+1}/n)) \leq \|v_{i+1} - v_i\|_{\mathcal{M}(z+v_i)} \leq \|v_{i+1} - v_i\| \sqrt{\|\mathcal{M}(z + v_i)\|} \leq M_0 V_0,$$

which implies that $\mathrm{d}_n(z) - \mathrm{d}_n(z + v/n) \leq sM_0V_0/n \leq d^2 M_0 V_0^d/n = \Delta_0/n$. Exchanging the roles of z and $z + v/n$ we obtain likewise $\mathrm{d}_n(z + v/n) - \mathrm{d}_n(z) \leq \Delta_0/n$, which concludes the proof. $\square$

Let $\varphi : \mathbb{R}^d \to \Omega$ be the canonical surjection. We denote by $\mathrm{d}_{\mathrm{per}}$ the distance on Ω defined for all $z, z' \in \Omega$ by

$$\mathrm{d}_{\mathrm{per}}(z, z') := \min\{\|Z - Z'\|;\ Z \in \varphi^{-1}(z),\ Z' \in \varphi^{-1}(z')\}$$

Corollary 2.8. *For each $n \geq n_0$ and for all $z \in \Omega$ we define*

$$\mathrm{D}_n(x) := \min_{z \in \Omega_n} \mathrm{d}_n(z) + \Delta_0\ \mathrm{d}_{\mathrm{per}}(z, x), \tag{36}$$

where the constants n_0, Δ_0, are defined in Lemma 2.7.

The map $\mathrm{D}_n : (\Omega, \mathrm{d}_{\mathrm{per}}) \to \mathbb{R}_+$ is Δ_0-Lipschitz, and $\mathrm{d}_n(x) = \mathrm{D}_n(x)$ for all $x \in \Omega_n$.

Proof. The map D_n is Δ_0-Lipschitz since it is defined as the minimum of the family of Δ_0-Lipschitz functions $x \mapsto \mathrm{d}_n(z) + \Delta_0\, \mathrm{d}_{\mathrm{per}}(z, x)$, indexed by $x \in \Omega_n$.

Let $x \in \Omega_n$, let $v = (v_1, \cdots, v_d) \in \mathbb{Z}^d$ and let $v_{\max} := \max\{|v_1|, \cdots, |v_d|\}$. Applying $v_{\max}$ times (35) we obtain

$$|\,\mathrm{d}_n(x) - \mathrm{d}_n(x + v/n)| \leq \Delta_0 v_{\max}/n \leq \Delta_0 \|v/n\|.$$

It follows that $\mathrm{d}_n(x) \leq \mathrm{d}_n(z) + \Delta_0\, \mathrm{d}_{\mathrm{per}}(x, z)$ for all $x, z \in \Omega_n$. This immediately implies as announced that $\mathrm{d}_n(x) = \mathrm{D}_n(x)$ for all $x \in \Omega_n$, which concludes the proof. $\square$

Before stating the main result of this section, we recall the definition of the viscosity solution of an eikonal equation [11].

Definition 2.9. *The viscosity solution of the eikonal equation (2), is the unique continuous function $\mathrm{D} : \Omega \to \mathbb{R}_+$ such that for any $\varphi \in C^1(\Omega, \mathbb{R})$ and any $z \in \Omega \setminus \{0\}$ the following holds:*

- *If $\mathrm{D} - \varphi$ has a unique global maximum at z, then $\|\nabla\varphi(z)\|_{\mathcal{M}(z)^{-1}} \leq 1$.*
- *If $\mathrm{D} - \varphi$ has a unique global minimum at z, then $\|\nabla\varphi(z)\|_{\mathcal{M}(z)^{-1}} \geq 1$.*

In alternative definitions of the notion of viscosity solution, the assumption that "$\mathrm{D} - \varphi$ has a unique global maximum at z" is often replaced with "$\mathrm{D} - \varphi$ attains a local maximum at z" (resp. minimum). These two definitions are equivalent, since one may subtract (resp. add) to φ a suitable smooth function $\psi \in C^1(\Omega, \mathbb{R}_+)$, large far from z, and with a parabolic behavior close to z: $\psi(z + h) \approx \lambda\|h\|^2$ for h sufficiently small.

We show in the following that the functions D_n, introduced in Lemma 2.8, converge uniformly towards the viscosity solution of (2) as $n \to \infty$. This immediately implies the convergence of the discrete maps d_n produced by our modified algorithm, as stated in Point (iii) of Theorem 2.2. The proof is similar in essence to the proof provided in [5], yet with a number of minor modifications due to our specific context.

Since $(\Omega, \mathrm{d}_{\mathrm{per}})$ is a compact metric space, since $\mathrm{D}_n(0) = 0$ and since D_n is Δ_0-Lipschitz for all $n \geq n_0$, the theorem of Arzelà-Ascoli implies that the sequence D_n is pre-compact. In the rest of the proof we consider an arbitrary subsequence $(\mathrm{D}_{\sigma(n)})_{n\geq 0}$ which converges uniformly on Ω towards a limit $\mathrm{D} : \Omega \to \mathbb{R}_+$. Our objective is to establish that D is a viscosity solution of the eikonal equation (2). Once this point is established, we conclude as announced that D_n converges uniformly toward the viscosity solution of (2), using the uniqueness of the viscosity solution and the pre-compactness of the sequence D_n.

We consider for each $n \geq 1$ a point $z_n \in \Omega_n$, the non-zero vertices $v_1^n, \cdots, v_d^n \in \mathbb{Z}^d$ of a common simplex $T \in \mathcal{T}(z_n)$, and some coefficients $\alpha_1^n, \cdots, \alpha_d^n \in \mathbb{R}_+$, $\alpha_1^n + \cdots + \alpha_d^n = 1$, which will all be specified later. We consider $\lambda_n > 0$, and $v_n \in \mathbb{R}^d$, $\|v_n\| = 1$, such that

$$\lambda_n v_n = \alpha_1^n v_1^n + \cdots + \alpha_d^n v_d^n.$$

Note that $\lambda_n \leq V_0$. On the other hand using the acuteness property (13) and denoting $M := \mathcal{M}(z_n)$ we obtain

$$\begin{aligned}
\lambda_n^2 \|M\| &\geq \|\lambda_n v_n\|_M^2 \\
&= \sum_{1\leq i\leq n} (\alpha_i^n)^2 \|v_i^n\|_M^2 + 2 \sum_{1\leq i<j\leq n} \alpha_i^n \alpha_j^n (v_i^n)^{\mathrm{T}} M v_j^n \\
&\geq \left(\sum_{1\leq i\leq n} (\alpha_i^n)^2\right) \min_{1\leq i\leq n} \|v_i^n\|_M^2 \\
&\geq d^{-1} \|M^{-1}\|^{-1},
\end{aligned}$$

hence $\lambda_n \geq (d\kappa(\mathcal{M}))^{-\frac{1}{2}}$. Denoting

$$D_n := \mathrm{d}_n(z_n) - \sum_{1\leq i\leq d} \alpha_i^n \,\mathrm{d}_n(z_n + v_i^n/n),$$

it follows from the definition (27) of the Hopf-Lax update that

$$D_n \leq \lambda_n \|v_n\|_{\mathcal{M}(z_n)}/n \tag{37}$$

but also that, given $n > 0$ and $z_n \in \Omega_n$, we may choose $v_1^n, \cdots, v_d^n$ and $\alpha_1^n, \cdots, \alpha_d^n$ in such way that the above inequality is an equality.

Consider an arbitrary map $\varphi \in C^1(\Omega, \mathbb{R})$ and a point $z \in \Omega \setminus \{0\}$. Denoting $L := \nabla\varphi(z)^{\mathrm{T}}$ we have for any $x \in \Omega$ and any $h \in \mathbb{R}^d$ the Taylor development

$$|\varphi(x+h) - \varphi(x) - Lh| \leq \int_0^1 \|\nabla\varphi(x+th) - \nabla\varphi(z)\| \|h\| dt \leq \omega(\mathrm{d}_{\mathrm{per}}(x, z) + \|h\|) \|h\|, \tag{38}$$

where ω denotes the modulus of continuity of the continuous function $\nabla\varphi : \Omega \to \mathbb{R}^d$. We denote

$$S_n := \varphi(z_n) - \sum_{1\leq i\leq d} \alpha_i^n \varphi(z_n + v_i^n/n),$$

and we observe that, using (38) and denoting $r_n := \mathrm{d}_{\mathrm{per}}(z, z_n) + V_0/n$,

$$|S_n + \lambda_n L v_n/n| \leq \sum_{1\leq i\leq n} \alpha_i^n \,|\varphi(z_n + v_i^n/n) - \varphi(z_n) - L v_i^n/n| \leq \omega(r_n) V_0/n. \tag{39}$$

We now explicit the choice of z_n, (v_i^n) and (α_i^n), and we use the above inequalities to conclude the proof. Let us assume that $\mathrm{D}-\varphi$ has a strict global maximum (resp. minimum) at z. For each $n>0$ we denote by $z_n \in \Omega_n$ a maximizer (resp. minimizer) of $\mathrm{D}-\mathrm{d}_n$ on Ω_n. It follows from the uniform convergence $\mathrm{D}_{\sigma(n)} \to \mathrm{D}$, as $n \to \infty$, that $z_{\sigma(n)} \to z$.

We consider a fixed $v \in \mathbb{R}^d$ such that $\|v\|=1$, and we choose $\alpha_1^n, \cdots, \alpha_d^n$ and $v_1^n, \cdots, v_d^n$ as above in such way that $v_n = v$ for all $n>0$. (resp. We choose $\alpha_1^n, \cdots, \alpha_d^n$ and $v_1^n, \cdots, v_d^n$ in such way that (37) is an equality for all $n>0$, and we denote by v an arbitrary cluster value of the sequence $(v_{\sigma(n)})_{n\geq 2}$.)

Using successively (39), the definition of z_n, and (37) we obtain

$$-\lambda_n L v_n/n - \omega(r_n) V_0/n \leq S_n \leq D_n \leq \lambda_n \|v_n\|_{\mathcal{M}(z_n)}/n$$

(resp. $-\lambda_n L v_n/n + \omega(r_n) V_0/n \geq S_n \geq D_n = \lambda_n \|v_n\|_{\mathcal{M}(z_n)}/n$). It follows that

$$-L v_n \leq \|v_n\|_{\mathcal{M}(z_n)} + \omega(r_n) V_0/\lambda_n$$

(resp. $-L v_n \geq \|v_n\|_{\mathcal{M}(z_n)} - \omega(r_n) V_0/\lambda_n$). Since $r_{\sigma(n)} \to 0$ as $n \to \infty$ (resp. $r_{\sigma(n)} \to 0$ and v is a cluster value of the sequence $(v_{\sigma(n)})_{n>0}$), and since λ_n is bounded below independently of n, we obtain

$$-Lv \leq \|v\|_{\mathcal{M}(z)}$$

(resp. $-Lv \geq \|v\|_{\mathcal{M}(z)}$). Observing that

$$\|\nabla\varphi(z)\|_{\mathcal{M}(z)^{-1}} = \sup\left\{\frac{-Lv}{\|v\|_{\mathcal{M}(z)}};\ v \in \mathbb{R}^d,\ \|v\|=1\right\}$$

we obtain as announced that $\|\nabla\varphi(z)\|_{\mathcal{M}(z)} \leq 1$ (resp. $\|\nabla\varphi(z)\|_{\mathcal{M}(z)} \geq 1$). It follows that D is the viscosity solution of (2), which concludes the proof.

3 Numerical experiments

This section is devoted to the numerical comparison of three solvers of the eikonal equation: two popular methods (AGSI, FM-8) which enjoy a reputation of simplicity and efficiency in applications, and the proposed algorithm (FM-LBR). The Adaptive Gauss Seidel Iteration[2] (AGSI) [5] produces numerical approximations which are guaranteed to converge towards the solution of the continuous anisotropic eikonal equation as one refines the computation grid[3], for an arbitrary continuous Riemannian metric $\mathcal{M}$. Fast Marching using the 8 point stencil (FM-8, stencil illustrated on Figure 1, center left) does not offer this convergence guarantee, but has a quasi-linear complexity $\mathcal{O}(N \ln N)$, in contrast to the super-linear complexity $\mathcal{O}(\mu(\mathcal{M}) N^{1+\frac{1}{d}})$ of the AGSI. Fast Marching using Lattice Basis Reduction[4] (FM-LBR) aims to offer the best of both worlds: a convergence guarantee, and fast computation times[5].

[2] As suggested in [5], the stopping criterion tolerance for the iterations of the AGSI is set to 10^{-8}.

[3] The grid is triangulated with a trivial periodic mesh, for the AGSI and the MAOUM

[4] We use for each discrete point z the $\mathcal{M}(z)$-reduced neighborhood described by Proposition 1.9 in 2-d, and Proposition 1.10 in 3-d, except if the matrix $\mathcal{M}(z)$ is detected to be exactly diagonal. In that case we use the standard 4 vertices neighborhood in 2-d (resp. 6 vertices in 3-d), which is a $\mathcal{M}(z)$-reduced mesh, see Figure 1 (left and center right). This modification has little impact on accuracy or CPU time, but avoids to pointlessly break the symmetry of the numerical scheme. A C++ source code, provided as an ancillary file to the Arxiv version of this paper, allows to reproduce the above experiments.

[5] Note that memory requirements are doubled for the FM-LBR in comparison with the AGSI and FM-8, see Remark 2.5.

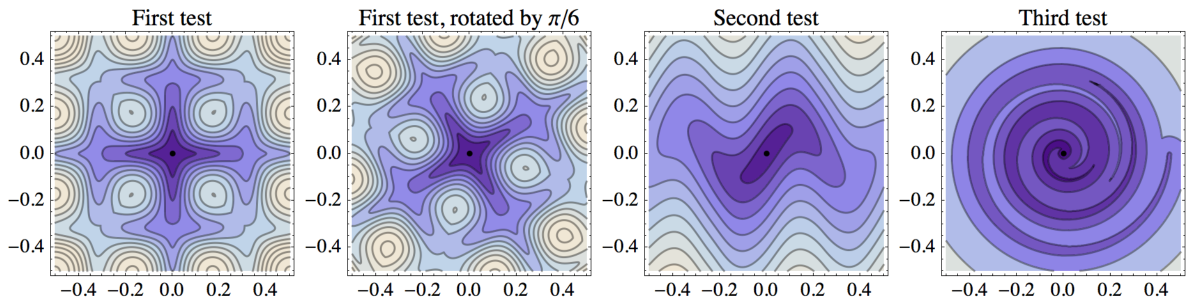

Figure 6: Level lines of the solutions of the two dimensional test cases.

We also implemented the Monotone Acceptance Ordered Upwind Method (MAOUM) [2], a Dijkstra inspired method using static stencils, like the FM-LBR. The difference between these two methods is that the stencil $V_n(z)$ used by the MAOUM at a point grid point $z \in \Omega_n$ is isotropic, only depends on the anisotropy ratio $\kappa(\mathcal{M}(z))$, and its boundary has cardinality $\#(\Omega_n \cap \partial V_n(z)) = \mathcal{O}(\kappa(\mathcal{M}(z)))$; in contrast the stencil of the FM-LBR is anisotropic, aligned with the ellipse defined by $\mathcal{M}(z)$, and of cardinality $\mathcal{O}(1)$. The stencils of the MAOUM were precomputed and stored in a look-up table, resulting in a complexity $\mathcal{O}(\kappa(\mathcal{M})N \ln N)$ for this algorithm.

We consider four test cases. The first two involve two dimensional Riemannian metrics $\mathcal{M}$ of limited anisotropy, $\kappa(\mathcal{M}) \simeq 5$, and were proposed in [24, 19] (including the grid sizes). All four methods (FM-LBR, FM-8, AGSI, MAOUM) are viable, yet the proposed method shows a reduced CPU usage and a competitive accuracy in comparison with its alternatives. The third test case, introduced in [4] and related to image segmentation problems, involves a highly anisotropic two dimensional metric $\mathcal{M}$, $\kappa(\mathcal{M}) \simeq 100$. It was observed in [4] that neither the AGSI or the FM-8 are viable in this context, due to complexity and accuracy issues respectively; and the MAOUM is no better. We show that FM-LBR solves this dilemma. The fourth test is a generalization of the third one in three dimensions.

Remark 3.1. *Contrary to the assumptions of the convergence analysis, in §2, but consistently with the envisioned applications, our test cases do not involve* periodic *boundary conditions. As a result, in exceptional circumstances, a few grid points usually in the domain's corners may not be reached by the FM-LBR, because they are not connected to the origin in the underlying graph (Lemma 2.7 does not apply). In the following experiments this only happened when the first test case was rotated by an angle $\theta \in [0.56, 0.61]$ radians. The numerical values at the four corners of the grid, equal to $+\infty$, were then rejected when computing the L^1 and L^∞ errors.*

The last two test cases involve discontinuous *Riemannian metrics, which happens in many applications (e.g. at the junction between materials of different index in geometrical optics), but also contradicts the assumptions of our convergence analysis.*

The first benchmark, introduced in [24], consists in computing the Riemannian distance from the origin $(0, 0, 0)$ on the surface parametrized by $f : (x, y) \to (x,\, y,\, (3/4)\sin(3\pi x)\sin(3\pi y))$. Denoting by $F := (\partial_x f, \partial_y f)$ the differential of f, the Riemannian metric is given by $\mathcal{M} = F^{\mathrm{T}} F$. The coordinates (x, y) are restricted to the square $[-0.5, 0.5]^2$, and the anisotropy ratio is $\kappa(\mathcal{M}) \simeq 5.1$ [24]. A reference solution is obtained on a 4000×4000 grid using the AGSI, and extended by bilinear interpolation. The FM-LBR, FM-8, AGSI and MAOUM are applied on a 292×292 grid, and the L^∞ and averaged L^1 errors estimated with respect to the reference

	FM-LBR	FM-8	AGSI	MAOUM
	First test			
CPU time	0.19	0.19	1.01	1.28
L^∞ error	3.99	1.47	1.62	8.80
L^1 error	1.13	0.53	0.51	2.33
	First test, rotated by $\pi/6$			
CPU time	0.20	0.21	1.44	1.31
L^∞ error	5.52	12.5	9.45	8.56
L^1 error	1.46	3.42	2.51	2.52
	Second test			
CPU time	0.076	0.079	0.77	0.36
L^∞ error	2.90	3.03	3.67	7.66
L^1 error	1.03	1.30	1.40	2.3

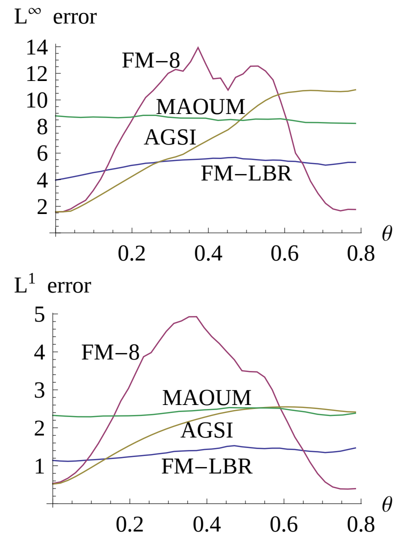

Figure 7: Tables of CPU time in seconds, L^∞ error an averaged L^1 error (left). Accuracy, in the first test rotated by an angle $\theta \in [0, \pi/4]$ (this interval is enough, since the dependence in θ is $\pi/2$-periodic and even). In average over theta, CPU times are $0.21s$, $0.20s$, $1.37s$, $1.31s$, L^∞ errors 5.16, 7.64, 6.86, 8.57 and averaged L^1 errors 1.34, 2.58, 1.95, 2.40 for the FM-LBR, FM-8, AGSI and MAOUM respectively. *All errors are multiplied by* 100*, for better readability.*

solution. As shown on Figure 7, FM-LBR is the fastest in terms of CPU time[6], but is less accurate than the AGSI or the FM-8.

Rotating the this test case by the angle $\theta = \pi/6$, and conducting the same experiment, shows a different story: the numerical error increases strongly for the AGSI and the FM-8, while the FM-LBR, unaffected, is now the most accurate method, see Figure 7. Unsurprisingly the CPU time is almost unaffected for the one pass solvers (FM-LBR, FM-8, MAOUM), while it slightly increases for the AGSI. The FM-LBR cuts L^∞ and L^1 numerical errors by 40% in comparison with the AGSI and the MAOUM, and CPU time by 85%, while the FM-8 produces even larger numerical errors.

Figure 7 shows that the FM-LBR offers the best accuracy for more grid orientations θ than its alternatives. The maximal error and averaged error with respect to θ are also in favor of the FM-LBR. The heuristic explanation of these empirical observations is the following: for $\theta = 0$, this test case is dominated by (close to) axis-aligned anisotropy. The fixed stencils of the AGSI and the FM-8 seem to benefit from this configuration; the FM-8 also works well for $\theta = \pi/4$, because its stencil includes the four diagonals. In contrast, the large stencils used by the MAOUM provide a good angular resolution in all directions, hence the method is mostly unaffected by rotations of the discretization grid. The FM-LBR does not suffer either from off-axis anisotropy, thanks to its adaptive stencils which are aligned with the anisotropy directions. In the special case of constant Riemannian metrics, random grid orientations are actually more favorable in average to the FM-LBR than axis aligned ones, see [12].

The second benchmark, discussed in [24, 19], is inspired by seismic imaging. The Riemannian metric $\mathcal{M}$ is defined as follows: at each point $z = (x, y)$, the symmetric matrix $\mathcal{M}(z)$ has

[6] All timings obtained on a 2.4Ghz Core 2 Duo, using a single core. Timings of the FM-LBR include the stencil construction, which typically accounts for 25%.

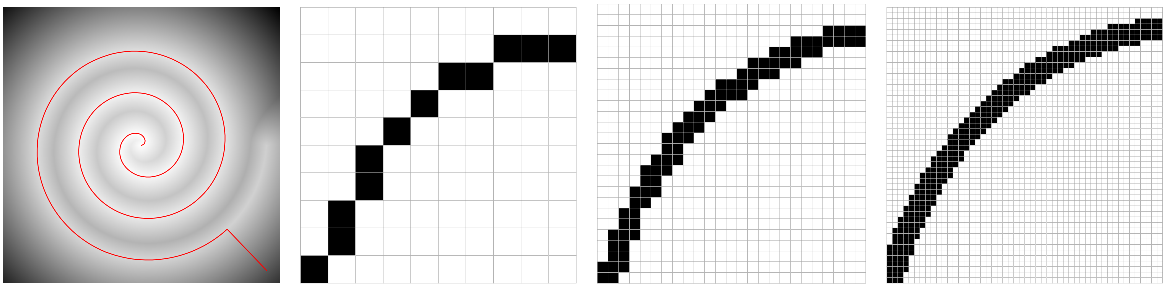

Figure 8: Reference solution for the third test case (left). The Riemannian metric $\mathcal{M}$ is anisotropic only on a thin band along a spiraling curve, wide of a few grid points. Detail at resolutions $n \times n$, where n equals 200 (center left), 500 (center right) and 1000 (right).

the two eigenvalues 0.8, 0.2, the latter associated to the eigenvector $(1, (\pi/2)\cos(4\pi x))$. There is little or no bias here towards axis-aligned anisotropy. The Riemannian distance from the origin $(0,0)$ is computed on the square domain $[0.5, 0.5]^2$. A reference solution is obtained on a 4000×4000 grid using the AGSI, and compared with the results of the four algorithms, obtained on a 193×193 grid. As shown on the table Figure 7, FM-LBR takes a smaller CPU time and offers a better accuracy than its alternatives.

The third [4] and fourth test cases are relevant benchmarks if one's objective is to use fast marching methods for the segmentation of tubular structures, in medical image or volume data respectively. The Riemannian metric $\mathcal{M}$ is euclidean (equal to Id) on a box domain Ω, except on the neighborhood of a curve Γ, see Figures 8 and 12 (right). The matrix $\mathcal{M}(z)$, when $z \in \Omega$ is close to Γ, has two eigenspaces: one of dimension 1, directed "tangentially" to the curve Γ, and associated to a small eigenvalue ($1/100^2$ and $1/50^2$ respectively), and one of co-dimension 1, associated to the eigenvalue 1. The (approximated) Riemannian distance D to to the origin $0 \in \Omega$ is computed, and a path of minimal minimal length γ joining a given point $P \in \Omega$ is extracted by "gradient descent on the Riemannian manifold $(\Omega, \mathcal{M})$":

$$\gamma'(t) = -\mathcal{M}(\gamma(t))^{-1} \nabla \mathrm{D}(\gamma(t)). \tag{40}$$

By construction of the Riemannian metric $\mathcal{M}$, traveling close and tangentially to the curve Γ is cheap. This is reflected by the level lines of D, and by the allure of the minimal path, see Figures 6, 8 and 12. Heuristically, this path joins the curve Γ in straight line, almost orthogonally, and then follows it. The alignment of the minimal path with the direction of anisotropy, observed in this test case, is not an uncommon phenomenon. The FM-LBR presumably benefits a lot from this behavior in terms of accuracy, since its stencils typically provide a good "angular resolution" in the direction of anisotropy, see Figures 2, 3, 5.

The performance of the FM-LBR, FM-8 and AGSI is illustrated on Figure 9 for the third test case. The MAOUM showed a poor accuracy in this test, presumably due to the huge stencils it generated, hence its results are not shown. The CPU time/resolution curve of the AGSI shows a stronger slope than the one pass solvers (FM-LBR, FM-8), which reflects its intrinsically larger complexity, namely $\mathcal{O}(\mu(\mathcal{M}) N^{1+\frac{1}{d}})$ instead of $\mathcal{O}(N \ln N)$. The L^∞ and L^1 error curves suggest that the FM-8 is not consistent in this test case, contrary to the FM-LBR and AGSI. The reference solution was obtained on a 4000×4000 grid, using the variant of the Fast Marching algorithm described in [13]. The AGSI could not be used due to prohibitive computation

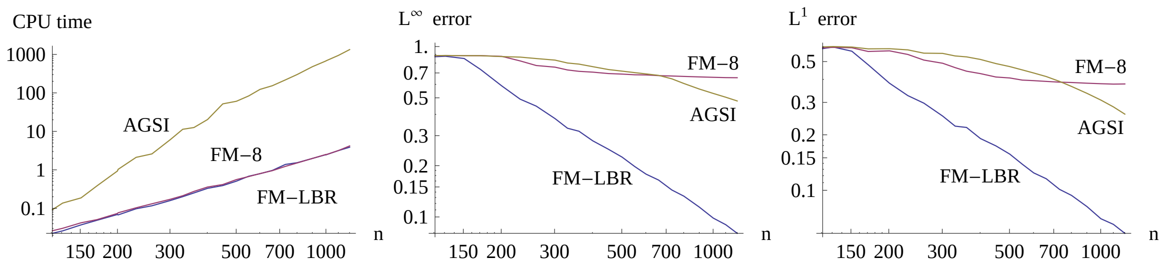

Figure 9: CPU Time (left, in seconds), L^∞ error (center), and averaged L^1 error (right) of the FM-LBR, FM-8 and AGSI, at several resolutions ranging from 120 to 1200 (log-log scale).

times and unimpressive accuracy in this test, and the FM-LBR was rejected because using it for reference might have induced a bias in our favor. At the resolution 1000×1000, typical in image analysis, the FM-LBR cuts the L^∞ error by 80% and the L^1 error by 75% with respect to the AGSI, while reducing the CPU time from 11 minutes to 2.5 seconds.

Obtaining the shortest path joining two given points is essential in motion planning control problems [2], as well as in the envisioned application [4]. This involves solving the Ordinary Differential Equation (ODE) (40), a task less trivial than it seems. The author is conscious that a benchmark based on minimal paths may reflect the properties of the ODE solver (and the time spent adjusting its sometimes numerous parameters), as much as those of the eikonal solver, but does so nevertheless due to the importance of minimal paths in applications. Eikonal solvers based on the discrete fixed point problem (10), such as the FM-LBR, FM-8 and AGSI, provide at each grid point $z \in \Omega_n$ an estimate $\mathrm{d}(z)$ of the distance $\mathrm{D}(z)$, and in addition an estimate $v(z)$ of the direction and orientation of the distorted gradient $-\mathcal{M}(z)^{-1}\nabla \mathrm{D}(z)$. This estimate has the form

$$v(z) := \frac{1}{n} \sum_{1 \leq i \leq k} \alpha_i v_i = \sum_{1 \leq i \leq k} \alpha_i (z_i - z), \tag{41}$$

where the integer $1 \leq k \leq d$, the positive barycentric coefficients $(\alpha_i)_{i=1}^k$ and the vertices $(v_i)_{i=1}^k$ of the mesh $\mathcal{T}(z)$ (resp. vertices $(z_i)_{i=1}^k$ of $V_n(z)$) achieve the minimum in the Hopf-Lax update operator (27).

From this point, a typical approach to solve (40) is to extend the values of d or v to the continuous domain Ω via an interpolation procedure, and then to use a black box ODE solver or a Runge Kutta method. Note that the accuracy usually expected from these high order methods is mostly doomed, since the discretization (10) of the eikonal equation is only first order, and since the vector field $\mathcal{M}^{-1}\nabla \mathrm{D}$ is discontinuous both at "caustics" and discontinuities of $\mathcal{M}$. A more significant issue is that computations frequently get stuck, despite the use of state of the art and/or commercial interpolation methods and ODE solvers, see e.g. [2] Figure 5.10.

We propose a method for the computation of minimal paths, which trades high order accuracy for robustness, and never gets stuck if the eikonal solver is Dijkstra inspired. It takes advantage of the specific form (10) of the discretization of the eikonal equation, and does not rely on black box routines. It is parameter free: there is not interpolation order or gradient step to adjust.

As illustrated on Figure 11, the FM-LBR recovers a qualitative minimal path in the third test case for grid resolutions much coarser than the AGSI. See Figure 8 for an illustration of the impact of coarse resolutions on the problem discretization. Hence the better accuracy of

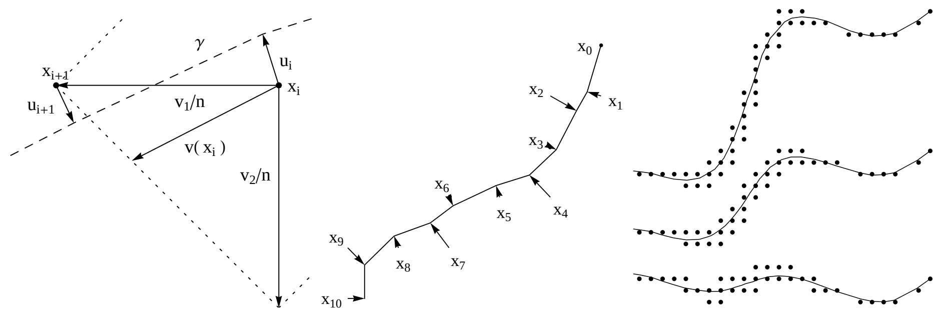

Figure 10: Notations for the minimal path computation (left); the contour of the stencil $V_n(x_i)$ is shown dotted. Grid points $x_0, \cdots, x_r \in \Omega_n$, corrections $u_0, \cdots, u_r \in \mathbb{R}^d$ shown as arrows, and piecewise linear path γ (center). Grid Points $(x_i)_{i=1}^r$ and piecewise linear path γ in the second test case at resolution $n \times n$, $n = 100$, using the FM-LBR (center right, detail).

the FM-LBR in this test case, see Figure 9, does effectively translate into a better extraction of minimal paths.

Our method produces a piecewise linear path γ, joining the points $x_0 + u_0, \cdots, x_r + u_r$ of Ω, where $x_0, \cdots, x_r \in \Omega_n$ are grid points, and $u_0, \cdots, u_r \in \mathbb{R}^d$ are small correcting offsets, see Figure 10. The first grid point x_0 is given by the user (and should satisfy $\mathrm{d}(x_0) < \infty$, see Remark 3.1), and the first offset u_0 is zero. All produced grid points x_{i+1}, $0 \leq i < r$, are vertices of the stencil $V_n(x_i)$ of the previous point. In the case of a Fast Marching method, the causality property implies that the values $(\mathrm{d}(x_i))_{i=0}^r$ are strictly decreasing, down to zero: $\mathrm{d}(x_r) = 0$ (the last point belongs to the source).

Minimal path computation, starting from a given grid point $P \in \Omega_n$.
Initialisation: $x_0 := P$, $u_0 := 0$, $i ::= 0$.
While $\mathrm{d}(z_i) > 0$ **do**

- Denote by $z_1, \cdots, z_k$ the grid points appearing in the expression (41) of $v(x_i)$.
- Find $\lambda \in \mathbb{R}_+$, $1 \leq j \leq k$, which minimize $\|x_i + u_i + \lambda v(x_i) - z_j\|$
- Set $x_{i+1} := z_j$ and $u_{i+1} := x_i + u_i + \lambda v(x_i) - z_j$.
- $i ::= i + 1$.

The piecewise linear path $\gamma : [0, r] \to \Omega$, parametrized so that $\gamma(i) = x_i + u_i$, satisfies the differential equation

$$\gamma'(t) = \lambda_{\lfloor t \rfloor} v(x_{\lfloor t \rfloor}), \tag{42}$$

for non-integer $t \in [0, r]$, where the constants $(\lambda_i)_{i=1}^r$ are the minimizers in the second step of the while loop, and the vector $v(z)$, $z \in \Omega_n$, is defined in (41). The particularity of our path extraction method is that the direction field v is not evaluated on the curve γ, but at the nearby points $(x_i)_{i=0}^r$. The next proposition shows that these points remain at distance $\mathcal{O}(\kappa(\mathcal{M})^3/n)$ from the curve in dimension 2 (the exponent 3 seems to be either over-estimated or a rare worst case scenario, in view of the experiments, see Figure 10), hence the accuracy of our path extraction method should be no worse than an explicit Euler integration of the ODE (40), with step $\kappa(\mathcal{M})^3/n$. We only use the fact that the meshes $\mathcal{T}(z)$, $z \in \Omega$, have diameter $\mathcal{O}(\kappa(\mathcal{M}))$, and are built of triangles of area 1/2.

Proposition 3.2. *Consider a piecewise linear path $\gamma \in C^0([0, r], \Omega)$ extracted with the above algorithm and parametrized as in (42), in dimension 2, after the Execution of the FM-LBR on*

Ω_n *for a Riemannian metric* $\mathcal{M}$*. Then for all* $t \in [0, r]$*, one has for some absolute constant* C

$$\|\gamma(t) - x_{\lfloor t \rfloor}\| \le C\kappa(\mathcal{M})^3/n. \tag{43}$$

Proof. We consider a fixed integer $0 \le i < r$, and observe that for all $t \in [i, i+1[$ one has, since γ is linear on this interval and $\gamma(i) = x_i + u_i$

$$\|\gamma(t) - x_{\lfloor t \rfloor}\| \le \max\{\|\gamma(i+1) - x_i\|,\ \|\gamma(i) - x_i\|\} \le \max\{\|u_{i+1}\| + \|x_{i+1} - x_i\|,\ \|u_i\|\}.$$

By construction, there exists a non-zero vertex v_j of the mesh $\mathcal{T}(x_i)$, appearing in (41), such that $x_{i+1} - x_i = v_j/n$. Combining Propositions 1.6 and 1.9 we obtain that $\|v_j\| \le 2\kappa(\mathcal{M})$, hence $\|x_{i+1} - x_i\| \le 2\kappa(\mathcal{M})/n$.

Our next objective is to bound $\|u_i\|$ uniformly. In order to avoid notational clutter, we rescale the variables of interest by a factor n: we set $u := nu_i$, $u' := nu_{i+1}$, and $v := nv(x_i)$. By construction (second step of the while loop in the path computation), we have for any $\lambda \in \mathbb{R}_+$ and any $1 \le j \le k$

$$\|u'\| \le \|u + \lambda v - v_j\|, \tag{44}$$

where the integer $1 \le k \le 2$, and the vectors $(v_j)_{j=1}^k$ are those appearing in the expression (41) of v. If $k = 1$, then choosing $\lambda = 1$, $j = 1$, and observing that $v = v_1$, we obtain $\|u'\| \le \|u + 1 \times v - v_1\| = \|u_i\|$.

We next assume that $k = 2$, and observe that $v = \alpha_1 v_1 + \alpha_2 v_2$ for some $\alpha_1, \alpha_2 \in \mathbb{R}_+^*$ such that $\alpha_1 + \alpha_2 = 1$, see (41). We consider an arbitrary but fixed $\mu \in]1, \infty[$ (e.g. $\mu := \sqrt{2}$), and define $w_1 := v_1\mu - v_2/\mu$, $w_2 := v_2\mu - v_1/\mu$. Note that $\|w_j\| \le (\mu + \mu^{-1}) \max\{\|v_1\|, \|v_2\|\} \le C_0\kappa(\mathcal{M})$ with $C_0 = 2(\mu + \mu^{-1})$. The matrix A of lines w_1 and w_2 satisfies $|\det A| = (\mu^2 - \mu^{-2})|\det(v_1, v_2)| = \mu^2 - \mu^{-2}$, since $|\det(v_1, v_2)| = 1$, see Point (b) of Definition 1.1. On the other hand $\|A\| \le \sqrt{\|w_1\|^2 + \|w_2\|^2} \le \sqrt{2}C_0\kappa(M)$, thus $\|A^{-1}\|^{-1} = |\det A|/\|A\| \ge C_1/\kappa(M)$ where $C_1 := (\mu - \mu^{-1})/(\sqrt{2}C_0)$. Therefore there exists $1 \le j \le 2$ and $\varepsilon \in \{-1, 1\}$ such that

$$\varepsilon u^{\mathrm{T}} w_j \sqrt{2} \ge \sqrt{(u^{\mathrm{T}} w_j)^2 + (u^{\mathrm{T}} w_j)^2} = \|Au\| \ge \|A^{-1}\|^{-1}\|u\| \ge C_1\|u\|/\kappa(\mathcal{M}).$$

Assume, without loss of generality, that $j = 1$ and $\varepsilon = 1$ satisfy the above expression. Then choosing $\lambda = 1/(\alpha_1 + \mu^2\alpha_2)$ and $j = 1$ in (44) yields, with $\nu := \alpha_2\lambda$

$$\|u'\|^2 \le \|u - \nu w_1\|^2 = \|u_i\|^2 - 2\nu u^{\mathrm{T}} w_1 + 4\nu^2\|w_1\|^2 \le \|u\|^2 - 2\nu C_1\|u\|/\kappa(\mathcal{M}) + \nu^2 C_0^2\kappa(\mathcal{M})^2.$$

If $\|u\| \ge C_2\kappa(\mathcal{M})^3$, where $C_2 := C_0^2/(2C_1)$, then the above inequality shows that $\|u'\| \le \|u\|$, since $\nu \le 1$. If $\|u\|$ is below this bound, then choosing $\lambda = 0$ in (44) yields $\|u'\| \le \|u\| + \|v_1\| \le C_2\kappa(\mathcal{M})^3 + 2\kappa(\mathcal{M})$. By an immediate induction argument we thus obtain $\|u_i\| \le (C_2\kappa(\mathcal{M})^3 + 2\kappa(\mathcal{M}))/n$, for all $0 \le i \le r$, which concludes the proof. □

The fourth test case is a generalization of the third one to three dimensions. It also involves a spiraling curve Γ, surrounded by a thin tubular neighborhood where traveling tangentially to Γ is 50 times cheaper than in orthogonal directions, or anywhere else in the domain, see Remark 3.3 for details. CPU time was 105s for the FM-LBR, while the AGSI took 480s and failed to recover the minimal path presented on Figure 12 (center) (a straight line joining the two endpoints was obtained instead). We do not perform a detailed benchmark in this case, which is similar in essence to the third one. Let us simply point out that FM-LBR addresses here a large scale (more than 10 millions grid points), strongly anisotropic ($\kappa(\mathcal{M}) = 50$) three dimensional shortest path problem, with a good accuracy and within reasonable CPU time on standard laptop computer.

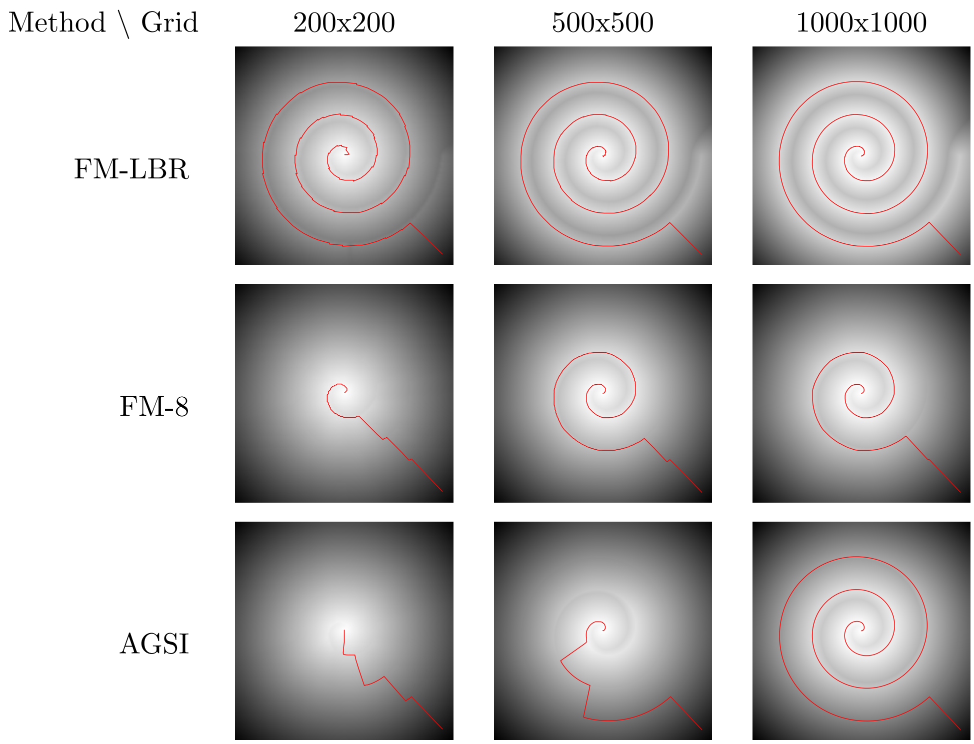

Figure 11: Visual comparison of the accuracy of three algorithms, at three resolutions, in the 2-d test case. Qualitatively, the approximate geodesic has the right behavior for a resolution as low as 170×170 with the FM-LBR, and 1000×1000 with the AGSI. This is presumably never the case for the FM-8, which is not consistent here.

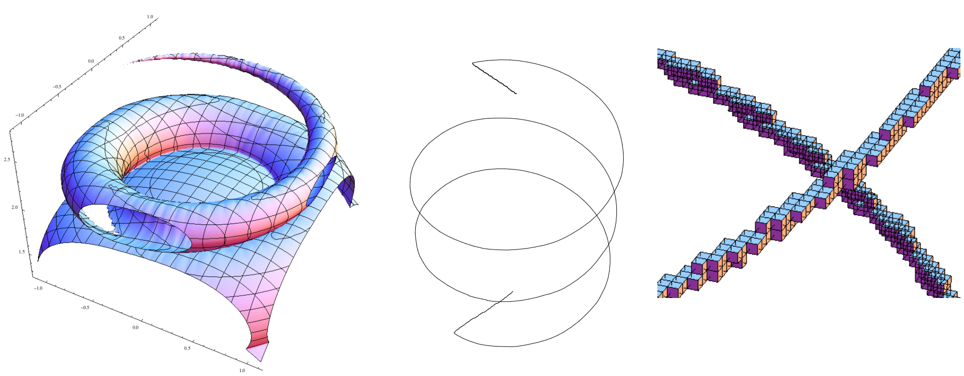

Figure 12: Results of the FM-LBR in the fourth, 3-d, test case. Iso-surface $\{\mathrm{d}(z) = 2\}$ (left), and shortest path joining the points $(0, 0, 0)$ and $(3, 0, 0)$ (center). Detail of the discrete points (represented by small cubes), in the neighborhood of the curve $\Gamma(t) = (\cos \omega_0 t, \sin \omega_0 t, t)$, for which the Riemannian metric is not euclidean (right).

Remark 3.3. *Exact description of third and fourth test cases, given for reference.*

- *The third, 2d test case is posed on the domain* $\Omega = [-1.1, 1.1]^2$*, involves three parameters* $\omega_0 = 6\pi$*,* $r_0 = \delta_0 = 1/100$ *(the parameter* δ_0 *was it seems slightly larger in the original experiment [4]), and the curve* $\Gamma : [0, 1] \to \Omega$ *defined by* $\Gamma(t) = t(\cos \omega_0 t, \sin \omega_0 t)$*. For all* $z \in \Omega$ *we have* $\mathcal{M}(z) = \mathrm{Id}$*, except if there exists* $0 \leq t \leq 1$ *and* $0 \leq r \leq r_0$ *such that* $z = \Gamma(t) + r(\cos \omega_0 t, \sin \omega_0 t)$*. In that case* $\mathcal{M}(z)$ *has the eigenvalues* δ_0^2 *and* 1*, the former being associated to the eigenvector* $\Gamma'(t)$*. Geodesic starting point:* $P = (1, -1)$*.*

- *The fourth, 3d test case is posed on the domain* $\Omega = [-1.1, 1.1]^2 \times [0, 3]$*, discretized on a* $200 \times 200 \times 272$ *grid, involves the three parameters* $\omega_0 = (5/2)\pi$*,* $r_0 = \delta_0 = 1/50$*, and the curve defined by* $\Gamma(t) = (\cos \omega_0 t, \sin \omega_0 t, t)$*. For all* $z \in \Omega$ *we have* $\mathcal{M}(z) = \mathrm{Id}$*, except if there exists* $t, \lambda, \mu \in \mathbb{R}$ *such that* $z = \Gamma(t) + (\lambda \cos(\omega_0 t), \lambda \sin \omega_0 t, \mu)$ *and* $\lambda^2 + \mu^2 \leq (r_0/2)^2$*. In that case* $\mathcal{M}(z)$ *has the eigenvalues* δ_0^2 *and* 1*, the former being associated to the eigenvector* $\Gamma'(t)$ *and the latter to the orthogonal space. Geodesic starting point:* $P = (0, 0, 3)$*.*

Conclusion

In this paper, we introduced a new discretization scheme for anisotropic eikonal equations: Fast Marching using Lattice Basis Reduction (FM-LBR). It is a variant of the classical Fast Marching algorithm, based on the algebraic concept of Lattice Basis Reduction, which strongpoints are the following. (I, Convergence) This algorithm is consistent for the anisotropic eikonal equation associated to any continuous Riemannian metric, of arbitrary anisotropy. (II, Complexity) It has a numerical cost comparable to classical isotropic Fast Marching, independently of the problem anisotropy. (III, Accuracy) The accuracy of the FM-LBR is competitive in general, and striking in test cases, related to tubular segmentation in medical images, where the Riemannian metric has a pronounced anisotropy close to and tangentially to a curve.

These strongpoints come at the price of the specialization of the FM-LBR: (i) the Riemannian metric may not be replaced with a more general Finsler metric, (ii) the domain needs to be discretized on a cartesian grid, and (iii) of dimension 2, 3, or in higher dimension the underlying Riemannian metric needs to have a block diagonal structure. Hopefully these requirements are met in many applications, and future work will be devoted to the application of the proposed algorithm in the context of medical image processing.

References

[1] K. Alton, I. M. Mitchell, *Fast Marching Methods for Stationary Hamilton-Jacobi Equations with Axis-Aligned Anisotropy*, SIAM Journal of Numerical Analysis, 47:1, pp. 363–385, 2008.

[2] K. Alton, I. M. Mitchell, *An Ordered Upwind Method with Precomputed Stencil and Monotone Node Acceptance for Solving Static Hamilton-Jacobi Equations*, Journal of Scientific Computing, 51:2, pp. 313–348, 2012.

[3] T. J. Barth, J. A. Sethian, *Numerical schemes for the Hamilton-Jacobi and level set equations on triangulated domains*, Journal of Computational Physics, 145(1), 1-40, 1998

[4] F. Benmansour, L. D. Cohen, *Tubular Structure Segmentation Based on Minimal Path Method and Anisotropic Enhancement*, International Journal of Computer Vision, 92(2), 192-210, 2010.

[5] F. Bornemann, C. Rasch, *Finite-element Discretization of Static Hamilton-Jacobi Equations based on a Local Variational Principle*, Computing and Visualization in Science, 9(2), 57-69, 2006.

[6] R. Gonzales, E. Rofman,*On Deterministic Control Problems: an Approximate Procedure for the Optimal Cost, I, the Stationary Problem*, SIAM Journal on Control and Optimization, 23, 2, pp. 242-266, 1985.

[7] S. Jbabdi, P. Bellec, R. Toro, J. Daunizeau, M. Pélégrini-Issac, H. Benali, *Accurate Anisotropic Fast Marching for Diffusion-Based Geodesic Tractography*, International Journal of Biomedical Imaging, 2008

[8] R. Kimmel, J. A. Sethian, *Computing geodesic paths on manifolds*, Proceedings of the National Academy of Sciences USA, 95(15), 8431-8435, 1998

[9] H.J. Kushner, P.G. Dupuis, *Numerical Methods for Stochastic Control Problems in Continuous Time*, Academic Press, New York, 1992.

[10] X.-G. Li, W. Yan, C. K. Chan, *Numerical schemes for Hamilton-Jacobi equations on unstructured meshes*, Numerische Mathematik, 94(2), 315-331, 2003

[11] P.L. Lions, *Generalized solutions of Hamilton-Jacobi equations*, Pitman, Boston, 1982.

[12] J.-M. Mirebeau, *On the Accuracy of Anisotropic Fast Marching*, to appear, preprint available on Arxiv, 2012.

[13] J.-M. Mirebeau, *Efficient Fast Marching with Finsler Metrics*, to appear, preprint available on Arxiv, 2012.

[14] P. Q. Nguyen, and D. Stehlé, *Low-dimensional lattice basis reduction revisited*, ACM Transactions on Algorithms, Article 46, 2009.

[15] G. Peyré, M. Péchaud, R. Keriven, L. D. Cohen, *Geodesic Methods in Computer Vision and Graphics*, Foundations and Trends in Computer Graphics and Vision, 5(3-4), 197-397, 2010.

[16] C. Rasch, T. Satzger, *Remarks on the $\mathcal{O}(N)$ Implementation of the Fast Marching Method*, IMA Journal of Numerical Analysis, 29, 806-813, 2009.

[17] M. Sermesant, E. Konukoglu, H. Delingette, *An anisotropic multi-front fast marching method for real-time simulation of cardiac electrophysiology*, Proc of Functional Imaging and Modeling of the Heart, 2007.

[18] J.A. Sethian, *Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision and Materials Sciences*, Cambridge University Press, 1996.

[19] J. A. Sethian, A. Vladimirsky, *Ordered Upwind Methods for Static Hamilton-Jacobi Equations: Theory and Algorithms*, SIAM Journal of Numerical Analysis, 41(1), 325-363, 2003

[20] J. A. Sethian, A. Vladimirsky, *Fast methods for the Eikonal and related Hamilton-Jacobi equations on unstructured meshes*, Proceedings of the National Academy of Sciences, 97(11), 5699-5703, 2000

[21] Y.-H.R. Tsai, L.-T. Cheng, S. Osher, and H.-K. Zhao, *Fast sweeping algorithms for a class of Hamilton-Jacobi equations*, SIAM Journal on Numerical Analysis, 41:2, pp.659-672, 2003.

[22] J. N. Tsitsiklis, *Efficient algorithms for globally optimal trajectories*, IEEE Transactions on Automatic Control, 40(9), 1528-1538, 1995.

[23] A. Vladimirsky, *Label-setting methods for Multimode Stochastic Shortest Path problems on graphs*, Mathematics of Operations Research 33(4), pp. 821-838, 2008.

[24] A. Vladimirsky, *Fast methods for static Hamilton-Jacobi Partial Differential Equations*, PhD Thesis, 2001

[25] A. Vladimirsky, *Static PDEs for Time-Dependent Control Problems*, Interfaces and Free Boundaries, 8(3), pp. 281-300, 2006.

[26] H. Zhao, *A Fast Sweeping Method for Eikonal Equations*, Mathematics of Computation, 74(250), 603-627, 2005